

(19)

Europäisches Patentamt  
European Patent Office  
Office européen des brevets



(11)

**EP 0 825 535 A2**

(12)

**EUROPEAN PATENT APPLICATION**

(43) Date of publication:  
25.02.1998 Bulletin 1998/09

(51) Int Cl.<sup>6</sup>: G06F 11/22

(21) Application number: 97306152.6

(22) Date of filing: 13.08.1997

(84) Designated Contracting States:  
AT BE CH DE DK ES FI FR GB GR IE IT LI LU MC  
NL PT SE  
Designated Extension States:  
AL LT LV RO SI

- Ellinger, Wesley M.  
Houston, Texas 77064 (US)
- Clark, Gordon R.  
The Woodlands, Texas 77380 (US)

(30) Priority: 16.08.1996 US 695828


(71) Applicant: Compaq Computer Corporation  
Houston Texas 77070 (US)

(74) Representative: Brunner, Michael John et al  
GILL JENNINGS & EVERY  
Broadgate House  
7 Eldon Street  
London EC2M 7LH (GB)

(72) Inventors:  
• Gready, Scott R.  
Spring, Texas 77379 (US)

(54) **A server controller for storing and retrieving sequences of video screens forwarded by a server host computer**

(57) A communication system is presented whereby sequences of video screens sent from a host CPU to a video controller can be stored and subsequently retrieved by a terminal located remote from the host CPU. The host CPU and video controller form part of a server arranged within a distributed computing system. An administrator situated at the remote terminal can retrieve select video screens produced during server operations to determine information regarding the server configuration and possible causes of server failure or future failure. The sequence of video screens thereby represent video screen changes stored upon a server controller adapted for coupling to the server expansion bus. The video screen changes represent a sequence of video screen changes occurring prior to server failure or after server reset. Those changes provide beneficial information to an administrator located remote from the server, and allows the administrator to communicate with the server using several possible communication protocols. The server controller snoops display data written from the host CPU to the video controller and mirrors the display data upon buffers within the server controller. Information within the buffers can be called upon by a remotely situated administrator regardless of whether server power is lost in the interim.

**EP 0 825 535 A2**

## Description

This invention relates to a distributed computing system, and more particularly to a host computer ("host") which can be accessed from a terminal not physically connected to the host and located at a site remote from the host. The remote terminal can access a stored sequence of video screens, such as a sequence of video screens occurring during host reset or failure operations. The sequence of video screens can then be replayed by a computer administrator located at the remote terminal. Remote access to those video screens allows the administrator to determine how a host operating system is responding to a reset, or possible reasons why the host system failed. Provided with the host is a printed circuit board ("PCB") which can be inserted into a backplane which accommodates the host. The PCB comprises a processor and memory for storing the sequence of video screens even when power is lost to the host.

Distributed computing systems are generally well known. Such systems allow communications between application programs hosted on numerous computer workstations. There are numerous types of distributed computing systems, often classified by the geographical extent of their communication capability. Terms used to classify the geographical breadth of distributed computing systems are, for example: local area networks (LANs), metropolitan area networks (MANs), and wide area networks (WANs).

Many of the more popular distributed computer systems employ a file server ("server"). Files, or data, are managed by a host within the server. Servers are particularly beneficial in allowing workstations fast access to files stored by the server. Accordingly, file servers embody a host computer which responds to an operating system program (a popular operating system being, e.g., Windows NT®) to not only orchestrate the files, but also to maintain file security, file backup, etc.

An important aspect of maintaining host functions within a server is to manage the host from a site remote from the host and, more specifically, to manage the host at a site remote from the workstations physically linked to the server. Recent trends have seen a steady increase in the number of servers used in business. Nowadays, servers are liberally used possibly at each location of a business entity -- rather than employing a centralized mainframe at one location. Unfortunately, funds available to administer many server hosts located at disparate locations is decreasing. While data placed on these servers is considered critical to the business, there remains insufficient means for ensuring their proper operation from a single service site. An expectation that an administrator travel to remote server sites to fix a problem is not only impractical but quite costly given the expense associated with server downtime.

Many operating systems, or applications associated with those operating systems, allow access to the host from a remote site often called a "virtual terminal". A virtual terminal, while not physically connected to the host, nonetheless allows remote control of certain operations of the host. Products such as COMPAQ Server Manager/R® ("SMR") and COMPAQ Insight Manager® ("CIM"), obtainable from Compaq Computer Corp., have attempted to address some of the issues involved in managing a network of distributed servers from a single, remote site. These products allow, *inter alia*, an administrator to be alerted as to a remote server failure, to reset the server from the remote site, and to access certain information provided on the server console.

It is certainly beneficial to allow remote control of certain server functions, especially those needed to reset one or more servers within a network of servers. Any downtime caused by server failure is probably the most costly time involved in running a distributed computer system. The causes of server failure, often termed server host "crash" are numerous. Any number of malfunctions or design flaws associated with the server hardware, server operating system or application programs running on the server may account for a server crash. If a server crashes, then file access is often lost and business records are temporarily inaccessible until the cause of failure is fixed.

A true benefit would result if an administrator located remote from the server can do more than be alerted to, and then reset, a failed server. In particular, it would be advantageous for the administrator to determine the cause of server failure so that he/she can possibly prevent future failures before they occur. Prevention of failure is as important, if not more important, than resetting a server that has crashed.

The cause of a failure is generally displayed on the server console at the time in which the server crashes. Moreover, irregularities in the server host hardware or operating system software can be detected upon reset (or "boot"). Those irregularities can lead to future failure if not attended to by the administrator. Accordingly, it would be beneficial to gain access to what is displayed on the server host console not only during server reset (or failure) but also leading up to server reset/failure. Information within the video screens (more particularly the sequence of video screens) displayed on the server console, which occur during server failure or reset would help remotely located administrators determine (and hopefully fix) an existing server failure or potential failure.

The video screens, resulting from a reset or failure of the server, comprise a sequence of video screen changes displayed on the host server console by the operating system, system basic input output system ("BIOS"), server application program or other system software. In particular, capture of two screen change sequences are of particular interest to a server administrator. In order to fix an existing failure or a future failure, it would be beneficial that the administrator be given the sequence of screen changes prior to server failure as well as the sequence of screen changes following a reset. Examples of server failure screens displayed on the server console are Microsoft Corp., Windows

NT® "blue screens" and Novell Corp., NETWARE® ABEND message which appear on the server console when the respective operating system crashes. These screens provide information such as processor fault indicia, system software routine addresses, and pertinent system memory contents. Upon reset of the server, the power on self test ("POST") code, associated with the aforementioned operating systems, typically performs some system diagnostics and displays information regarding failures detected to the server console screen. Hence, a means for capturing such sequences and replaying them at a remote management site is desired.

In addition to hardware and software problems, a server can also fail if power to the server is halted. Unfortunately, if power is halted, any screen changes as to what occurred prior to failure will be lost. A server is therefore needed which employs a mechanism for saving reset and failure screen changes even when power to the server is lost. The stored screen changes may then be beneficially read at a future date by a remotely situated administrator. The desired mechanism is one which can therefore maintain screen information during power loss, and can selectively forward power only to critical units within the mechanism. Accordingly, a mechanism is needed which is preferably embodied upon a PCB mountable within a server chassis. The PCB is desirably connected to the server host and includes media for storing screen information output from the host, and for maintaining that information even when server power is discontinued.

Communication between a remote site and a server is typically performed via text-based connection protocols, generally known in the industry as American National Standards Institute ("ANSI") terminal emulation protocols. Although terminal emulation protocols provide a certain level of functionality, it is desirable that other protocols, in particular protocols which enable application layer protocols such as simple network management protocol ("SNMP"), a protocol for communication of server management information, be supported on a point-to-point ("PPP") communications link between the server and the remote site. If a server is to include a PCB embodying a system for communicating with the remote site using a plurality of communications protocols, then it is desirable that subsystems upon the PCB determine which of the supported protocols (i.e., text-based, PPP, etc.) the remote site is using as it communicates with the server.

The problems outlined above are in large part solved by a remote communication system of the present invention. The remote communication system employs a server controller which can be connected to an expansion bus upon the server. Preferably, the server controller is embodied upon a PCB having edge connectors connectable to an EISA expansion bus. The server controller comprises detection logic which detects write cycles within a video address range. The write cycles are initiated as display data forwarded to a video controller also arranged upon the expansion bus. Detection logic thereby snoops the expansion bus for display data, and then mirrors the display data to a controller memory.

The controller memory, like the detection logic, is embodied upon the server controller PCB and includes a plurality of buffers. Display data forwarded to the controller memory is captured in a local frame buffer which stores frame-by-frame the display data. A previous frame or screen of display data is compared with a current frame or screen of display data to indicate a possible change. The change, as registered on the server display, is stored within current sequence buffers also associated with the controller memory. The current sequence buffers store three types of changes: a sequence of video screen changes which occur prior to server failure or reset, a sequence of video screen changes which occur after the most current server reset, and a sequence of video screen changes which occur after a reset which occurs prior to the most recent reset. Storing current and previous reset video screen changes as well as failure video screen changes (i.e., blue screens or ABEND messages) allows an administrator to determine reasons for server failure or possibilities of future failures.

A terminal located remote from the server can be used by an administrator to draw upon any of the various sequences of video screen changes stored within the controller memory. The server controller can recognize different types of communication protocols sent from the remote terminal. A local processor within the server controller can recognize the communication protocol being sent by employing state machines dedicated to each protocol. Depending upon a character or characters sent, the state machine can make a determination whether the communication protocol is a text-based protocol, a PPP protocol, or a pre-PPP protocol. The various protocols described immediately herein above are well known to the skilled artisan as protocols frequently used as inter-networking protocol used in performing out-of-band connection over, e.g., a phone line or direct serial connection. It is the ability to service many protocols from a remote terminal, and to detect the incoming protocol that defines an advantage hereof. Thus, regardless of the protocol sent, the processor is configured to act upon and forward the communication information to the controller memory. The controller memory can thereby dispatch various video screen sequences in response to directives embodied upon the communication protocol.

The detection logic, controller memory, processor and communication unit are compiled as a system configured upon a PCB. If power to the server should fail, then the primary power supply to the PCB via the expansion bus will also fail. The PCB is configured, however, to maintain certain vital functions even if power from the expansion bus ceases. The PCB employs a secondary power supply, or battery, which supplies power to the processor and controller memory, and selectively to the communication unit when expansion bus power ends. The battery will maintain processor

and controller memory activities as well as communication therebetween for a period of time sufficient to sustain video screen information stored within the controller memory. If the communication unit is not being accessed by a remote terminal, then power from the battery to the communication unit is terminated. Thus, decoupling circuitry is present which prevents unnecessary drain upon the battery if the communication unit is not being used. A first decoupling unit is present between a plane (or portion) of the PCB upon which the processor and the controller memory reside and a plane of the PCB upon which the communication unit resides. Another decoupling unit is present between the processor and controller memory plane and the plane of the PCB upon which the detection logic resides. The detection units thereby serve not only to decouple power between planes, but also to decouple signal conductors sent between planes. Decoupling power and tri-stating signals between planes helps preserve the charge upon the battery power source so that power needed to maintain the storage sequences of video screen changes during the server downtime is maximized.

Other objects and advantages of the invention will be apparent upon reading the following detailed description and upon reference to the accompanying drawings in which:

- 15 Fig. 1 is a block diagram of a host server retrofitted with a server controller accessible by a remote terminal, in accordance with the present invention;
- Figs. 2a-2c are three screen displays of menus which the server controller forwards to a display screen of the remote terminal in response to access from the remote terminal;
- 20 Fig. 3 illustrates portions of the memory address space and I/O address space available to the video memory and video control registers (within the video controller) shown in Fig. 1;
- Fig. 4 is a detailed block diagram of the server controller;
- Fig. 5 is a block diagram of various buffers used by the processor shown in Fig. 4 to process and store sequences of video screens sent to the expansion bus from the server CPU;
- 25 Figs. 6a and 6b are flowcharts illustrating steps taken by the server controller in storing certain sequences of video screens;
- Fig. 7 is a flowchart illustrating steps taken by the processor of Fig. 4 to replay a reset sequence of video screens on the remote terminal;
- Fig. 8 is a flowchart illustrating steps taken by the processor of Fig. 4 to replay a failure sequence of video screens on the remote terminal;
- 30 Figs. 9a-9c represent a sequence of displays presented as an exemplary sequence of failure screens;
- Fig. 10 is a timing diagram of various EISA expansion bus signals, indicating two separate video write cycles in which expansion bus intrusion and non-intrusion occur;
- Fig. 11 is a flowchart illustrating steps taken by the server controller to detect various communication protocols useable by the remote terminal;
- 35 Fig. 12 is a state diagram illustrate a text protocol state machine used in detecting text communication protocol sent from the remote terminal;
- Fig. 13 is a state diagram illustrate a pre point-to-point (PPP) protocol state machine used in detecting pre-PPP communication protocol sent from the remote terminal;
- Fig. 14 is a state diagram illustrate a PPP protocol state machine used in detecting PPP communication protocol sent from the remote terminal;
- 40 Fig. 15 is the server controller embodied upon a PCB adapted for coupling onto an expansion bus;
- Fig. 16 is comparator logic used in decoupling signal conductors extending between power planes arranged within said PCB of Fig. 15; and
- 45 Fig. 17 illustrates signals used to couple signal conductors and power to a communication unit power plane during times when the communication unit is receiving a signal from a remote terminal.

While the invention is susceptible to various modifications and alternative forms, specific embodiments thereof are shown by way of example in the drawings and will herein be described in detail.

- Referring now to Fig. 1, a block diagram of a computer system 10 is shown. Computer system 10 preferably includes any system which stores and retrieves files, a suitable system being a file server, an application server, or any other server known in a distributed computing system. System 10 includes a host central processing unit ("CPU") 12 coupled to a memory controller 14 via a CPU local bus. CPU 12 includes any data processing unit which implements a predetermined instruction set and has basic features of, *inter alia*, execution units, control/timing units and various registers. Exemplary CPU 12 includes the Pentium® processor manufactured by Intel Corp.
- 55 Memory controller 14 orchestrates the transfer of data between the CPU local bus and system memory 16. Memory controller 14 responds to various addressing techniques, and is configured to support various storage cell architectures within system memory 16, suitable architectures being dynamic random access memory ("DRAM") or static random access memory ("SRAM"). Memory controller 14 preferably operates synchronously with CPU 12 to ensure maximum

transfer bandwidth to and from system memory. According to one embodiment, CPU 12 is coupled to a peripheral component interface ("PCI") bus. A bus interface unit 18 operates as a bus bridge between the PCI bus and an expansion bus. Bus interface unit 18 includes a buffer to control the transfer of data and address signals between the PCI bus and the expansion bus. Bus interface unit 18 thereby interfaces between data sent over buses operating at possibly different clock speeds and with dissimilar protocols. Preferably, the expansion bus comprises an extended industry standard architecture ("EISA") bus configuration.

Connected to the expansion bus is a video controller 20. Video controller 20 is used to control an electronic display 24 connected thereto. Video controller 20 is indicative of commonly used video controllers such as video graphics array ("VGA") controllers. Video controller 20 responds to addresses upon the expansion bus and forwards corresponding data to video memory 22. Video memory 22 stores video data to be displayed by the video controller upon electronic display 24. Display 24 includes any console or monitor which can receive video data, a preferred display 24 is a cathode ray tube ("CRT") or, alternatively, a liquid crystal display ("LCD"). Video controller 20 and video memory 22 are henceforth referred to collectively as a video subsystem.

Server 10 is retrofitted with a mechanism for saving sequences of video screen changes prior to server 10 failure as well as sequences of video screen changes after server 10 reset. The mechanism for doing so must be unaffected by power loss to the server, and must have the ability to forward the stored video screens to a remote terminal. Such a mechanism is provided by a server controller 26. Server controller 26 is coupled within the chassis of server 10, and is connectable to the expansion bus, wherein controller 26 is accessible from a terminal located remote from server 10. According to a preferred embodiment, remote terminal 28 is coupled to server controller 26 via telephone line communication devices, such as a modem contained in both remote terminal 28 and server controller 26. Remote terminal 28 includes a display apparatus (e.g., a CRT) and an input device (e.g., a keyboard).

According to one embodiment, server controller 26 operates from software available from Compaq Computer Corp. Specifically, Compaq offers a software product entitled Compaq Insight Manager ("CIM") which can operate upon remote terminal 28. CIM can access firmware loaded into server controller 26 for performing, e.g., replay of the various video screen changes stored in server controller 26.

Hardware and firmware features of server controller 26 make available remote access and control by a remote terminal configured with CIM. Hardware and associated code added to server controller 26 allow a system administrator to dial into the server controller 26 after server 10 has failed, and thereafter communicate with the server controller 26 to view the sequences of display 24 screen changes that occurred prior to that failure and after reset.

The firmware or code within server controller 26 resides as console application code used to communicate with remote terminal 28. Fig. 2a is a screen display of the initial menu screen which the console application displays upon remote terminal 28. The screen displays occur whenever the system administrator dials into server controller 26. The administrator can choose from any of the menu items. If he or she chooses menu item four, then the administrator can view, at his or her remote terminal 28, the sequence of video screens stored within server controller 26. Three different sequences of video screen changes may be played: (i) a sequence of video screens which occur after a prior reset of server 10, (ii) a sequence of video screens which occur after a current reset of server 10, and (iii) a sequence of video screens which occur prior to a failure of server 10. The aforementioned sequence of video screens are herein referred to as current reset sequence, prior reset sequence and failure sequence (collectively "playback sequences").

Fig. 2b is a screen display of a menu screen which the system administrator can use to instruct server controller 26 to display any of the aforementioned sequences. Fig. 2c is a screen display of a menu screen which the system administrator uses to instruct the server controller 26 to start, pause, or change the speed of the playback of any of the current reset, prior reset and failure screen sequences.

The current reset screen sequence consists of a sequence of video screen changes that occur immediately after the most recent reset of server 10. The previous reset sequence consists of a sequence of video screen changes that occur immediately after a reset which occurred prior to the most recent reset of server 10. The failure sequence consists of a sequence of video screen changes that occur just prior to a reset, failure or power down of server 10. The number of screen changes which the server controller 26 stores prior to a reset or power down of the server 10 for a failure sequence or subsequent to a prior or current reset sequence is determined by the amount of buffer space available in the server controller 26 for the given playback sequence.

Referring now to Fig. 3, portions of the memory address space and I/O address space in server 10 are shown. The memory address range 0xB0000 through 0xBFFFF is reserved for video frame buffers which reside within video memory 22. The I/O address ranges 0x03B4 through 0x03BA and 0x03C0 through 0x03DA are designated for video control registers within video controller 20. CPU 12 writes values to the video control registers to control various aspects of display 24. The frame buffer memory address range and the video control register I/O address range are herein collectively referred to as video address ranges, or address ranges for data sent to the video subsystem.

Display 24, when in text mode, is organized into a matrix of rows and columns. For example, a common display screen configuration is 25 rows by 80 columns. Thus 2000 (i.e., 25 x 80) character locations are available on display 24. Each location on display 24 has an associated word (i.e., two bytes) in the frame buffer. The low byte of the word

contains the character value (such as a 0x41 for the character "A" in the ASCII character set) and the high byte of the word contains an associated attribute (such as the colour, foreground and background intensity, whether or not the character is blinking, etc.). CPU 12 writes values to the appropriate frame buffer locations within video memory 22 to display the desired characters and attributes on display 24. The video data in the frame buffer comprising characters and attributes for display is herein referred to as display data. In the above example, 4000 bytes (2000 locations x 2 bytes per location -- one for the character and one for the attribute) of the frame buffer are used to store display data representing an entire display screen.

The following three examples illustrate some of the many uses of the video control registers, whose addresses are reserved within the I/O address space. First, CPU 12 writes to the video control registers (i.e., within address ranges 0x03B4 to 0x03BA, or within address ranges 0x03C0 to 0x03DA) to control, e.g., the location and attributes of the cursor on display 24. Second, the CPU 12 writes to the video control registers to change, e.g., the video mode from colour mode to monochrome mode or vice versa. Third, the video control registers can specify a start address in video memory 22 from which multiple screens can be stored. In the third example, an assumption is made whereby video memory 22 contains more storage locations than are needed to store one screen of display data. Assuming that video memory 22 contains 64KB of memory (i.e., 64KB are available for the frame buffers), then possibly only 4000 bytes of the frame buffer are used to store a screen of display data. Thus, the frame buffer may accommodate multiple sets of display data at a time. CPU 12 writes to the video control registers to specify the start address in video memory 22 of the first set of display data to be displayed on display 24. Information regarding video controllers and more specifically VGA controller which employ video control registers for performing many of the above exemplary control functions is obtainable from, *inter alia*, Cirrus Logic Corp., part no. CL-GD542x.

Video controller 20 decodes the addresses generated on the expansion bus (preferably EISA bus) and responds to bus cycles in the video address ranges. Video controller 20 not only sends to display 24 the display data stored within the video memory 22, but also controls the display 24 in response to data written by CPU 12 into the video control registers. Video memory therefore contains character and attribute information, whereas video control registers contain information to control display 24.

Similar to video controller 20, server controller 26 also decodes the addresses on the expansion bus in order to detect writes to the video subsystem. When it detects a write to the video address ranges, server controller 26 stores the data written to the video address ranges within the server controller 26. This operation is referred to as "snooping". That is, server controller 26 obtains one or more screens of display data intended for another subsystem (i.e., video subsystem 20 and 22). Data is duplicated upon server controller 26 such that server controller 26 maintains a mirror copy of information within frame buffers of video memory 22 as well as information within video control registers of video controller 20. Snooping the expansion bus for digital signals within the video address range allows duplicative stores to occur, even though CPU 12 perceives only one write to the video subsystem. Importantly, snooping is performed by server controller 26 in a non-intrusive manner. Data intended for the video subsystem proceeds onto the video subsystem and is not modified or removed by server controller 26.

Referring now to Fig. 4, a detailed block diagram of server controller 26 is shown. Server controller 26 comprises detection logic 30 coupled to the expansion bus. Controller memory 32 is coupled to detection logic 30 via a controller bus. Detection logic 30 comprises combinatorial and sequential logic elements such as those commonly found, for example, in programmable array logic ("PAL") circuits. Details regarding control logic 40, address buffers 42 and data transceivers 44 will be provided after a discussion of their relevance. In particular, the three units within detection logic 30 gain significance given full appreciation to the capabilities of the various components of the server controller 26 and specifically the significance of controller memory 32 and the various buffers therein. Therefore, description of units 40, 42 and 44 will be forestalled until after a discussion of the playback sequences and how various buffers store those sequences. As part of server controller 26, a processor 34 is needed to orchestrate information upon the controller bus -- either information between detection logic 30 and controller memory 32, or information between communication unit 36 and controller memory 32. Processor 34 includes any microprocessor or microcontroller responsive to an instruction set and containing, e.g., an execution unit, memory, input/output devices, and one or more registers. Processor 34 is indicative of microprocessors used for embedded system applications. The controller memory 32 comprises a plurality of digital storage elements. In one embodiment the controller memory 32 comprises pseudo-static RAM. Communication unit 36 includes any device which can communicate to remote terminal 28 configured remote from server controller 26. Communication unit 36 responds to an analog ring signal (ARNG) sent from remote terminal 28, by producing a communication signal (COMM. SIG) back to remote terminal 28. Preferably communication unit 36 comprises a modem. Communication unit further contemplates a serial port coupled to processor 34 and a modem coupled to the serial port external to server controller 26 for communicating with the remote terminal 28.

Processor 34 stores the playback sequences in controller memory 32. Processor 34 could store the playback sequences as sequences of display data representing an entire set of display screens within controller memory 32. However, this would be a relatively inefficient way to use the controller memory 32 and thus the number of video screens stored within controller memory 32 comprising the playback sequences would be relatively small. In the present inven-

tion. processor 34 advantageously processes display data to efficiently store the playback sequences as a sequence of display "changes" instead of a sequence of entire display screens. The amount of display data representing changes within a screen is substantially less than the amount of display data representing an entire screen. Thus, the sequence of current reset screens, the sequence of prior reset screens and the sequence of failure screens (i.e., playback sequences) are represented as changes from a reference or "start" screen.

Referring now to Fig. 5, various buffers within controller memory 32 are shown. Those buffers are used by processor 34 to process and store the playback sequences. A local frame buffer 48 contains the display data snooped during write accesses to video memory 22. That is, detection logic 30 snoops the expansion bus and stores the data destined for the video subsystem into the local frame buffer 48. A snapshot buffer 50 contains a copy or snapshot, created by the processor 34, of the current display data from the local frame buffer 48. A temporary snapshot buffer 52 also contains a snapshot of the current display data from the local frame buffer 48. The temporary snapshot buffer 52 is used to make the playback sequences more legible in the event that screen changes, in particular screen scrolls, are rapidly made.

A current reset sequence buffer 54 contains information associated with the most recent reset sequence. That is, the current reset sequence buffer 54 contains the screen changes beginning with the most recent reset and ending when the current reset sequence buffer 54 becomes full. Screen changes which occur after the current reset sequence buffer 54 is full are ignored. The previous reset sequence buffer 56 contains information associated with the reset sequence just prior to the most recent reset sequence. That is, the previous reset sequence buffer 56 contains the screen changes beginning with a reset which occurred prior to the most current reset and ending when the previous reset sequence buffer 56 becomes full or when the next reset occurs. Screen changes which occur after the previous reset sequence buffer 56 is full are ignored.

A current sequence buffer 58 contains as many of the most recent screen changes that will fit within buffer 58. A failure sequence buffer 60 contains information associated with the failure sequence just prior to the most recent reset of server 10. That is, failure sequence buffer 60 contains the screen changes ending with the most recent reset and contains as many of the screen changes occurring just prior to the most recent reset that will fit therein.

The failure start buffer 64 contains display data of a video screen display associated with failure sequence buffer 60. The failure start buffer 64 contains the initial screen image (or "reference" display screen) in the failure playback sequence and consequently the screen image to which the first change in the failure sequence buffer 60 is applied. Initially the failure start buffer 64 is a blank screen. If and when the failure sequence buffer 60 overflows, the overflow changes are converted to characters and attributes and stored in the failure start buffer 64. Thus, the failure start buffer 64 provides a more useful initial screen image in the event where only a few characters on the screen changed many times over a relatively long period of time. The current start buffer 62 serves the same function relative to the current sequence buffer 58.

It is noted that a reset or failure of server 10 does not cause a reset of server controller 26. However, server controller 26 detects and maintains information concerning the following reset events: power on of server 10, power off of server 10, and reset of the expansion bus. In particular, server controller 26 does not respond to the RSTDRV EISA bus signal, but merely detects the assertion of RSTDRV. If a reset occurs, or if power is removed from server 10, server controller 26 nonetheless maintains operation in order to preserve the playback sequences for display on the remote terminal 28. A secondary power supply is furnished on the same PCB which houses server controller 26 to ensure the playback sequences are maintained. If desired, however, server controller 26 may be reset under software control by writing a predetermined sequence of values into predetermined register locations of the server controller 26 which are accessible to the CPU 12.

When a server 10 reset occurs, pointers are changed so that the current reset sequence buffer 54 becomes the previous reset sequence buffer 56, and previous reset sequence buffer 56 is flushed and becomes the current reset sequence buffer 54. A server reset also causes the current sequence buffer 58 to become the failure sequence buffer 60, and failure sequence buffer 60 is flushed and becomes the current sequence buffer 58. Also on reset, the current start buffer 62 becomes the failure start buffer 64 and the failure start buffer 64 is flushed and becomes the current start buffer 62.

Server controller 26 is configured to perform an automatic server recovery ("ASR") reset of server 10. When the server controller 26 detects a failure of server 10, it waits a predetermined amount of time (typically 30 seconds) and resets server 10.

The following is an exemplary sequence of a reset, followed by a failure, followed by a reset on server 10. Server 10 is first powered on which causes server 10 to reset. On reset, server controller 26 flushes information from previous reset sequence buffer 56, flushes failure sequence buffer 60 and flushes failure start buffer 64. During flushes of buffers 56, 60 and 64, a swap of previous reset sequence buffer 56 and the current reset sequence buffer 54 occurs simultaneous with swaps between buffers 60 and 58 as well as swaps between buffers 64 and 62. Thus, reset causes movement of information from buffers 54, 58 and 62 to buffers 56, 60 and 64, respectively. On reset, server 10 software begins performing the POST operations and displaying characters on display 24. Server controller 26 stores the screen

changes associated with the POST operations in current reset sequence buffer 54 and current sequence buffer 58. Once the current reset sequence buffer 54 becomes full, the server controller 26 stops storing screen changes thereto. If the current sequence buffer becomes full, the server controller 26 removes the oldest screen changes from the current sequence buffer 58 to make room for the new screen changes and modifies the current start buffer 62 to reflect the old screen changes. Thus, buffer 58 performs similar to a FIFO register, with overflow information sent to buffer 62. As described below, a combination of buffers 58 and 62 (when swapped with buffers 60 and 64) depict the entire failure sequence of video screens prior to the most recent failure.

When a failure of server 10 occurs, in accordance with the second step in the reset-failure-reset three-step example, a sequence of video screens occur upon display 24 indicative of that failure and possibly the cause of that failure. These screens are automatically displayed as part of many server operating systems available as ABEND messages from Novell Netware or a Microsoft NT blue screen. Thirty seconds after failure occurs, server controller 26 performs an ASR reset on server 10. Thus, the second reset within the reset-failure-reset example occurs during ASR reset. The reset, similar to the first reset, enables server controller 26 to flush buffers 56, 60 and 64 while swapping buffers 56, 60 and 64 with buffers 54, 58 and 62, respectively. Similar to the first reset, server 10 again initiates POST operations and forwards characters to display 24. Server controller 26 stores the screen changes occurring after reset, resulting from POST operation, in current reset sequence buffer 54 and current sequence buffer 58. Previous reset sequence buffer 56 now contains the sequence of screen changes associated with the initial power on reset, and the failure sequence buffer 60 as well as failure start buffer 64 now contain the sequence of screen changes leading up to that failure. As opposed to the initial reset contained within the previous reset sequence buffer 56, current reset sequence buffer 54 contains the reset associated with the second reset (or ASR reset).

Video screen changes, as opposed to the entire video screen, are stored as packets in the current reset sequence buffer 54, previous reset sequence buffer 56, current sequence buffer 58 and failure sequence buffer 60. Each packet in buffers 54, 56, 58 and 60 consists of a series of bytes of a specified length. Packets are classified according to a "type", and the packet type is determined by its length along with the first byte (byte 0), and second byte (byte 1) if the packet length is greater than one. The packet types, and the resulting action of server controller 26 upon remote terminal 28, are as follows:

Packet Type	Action
Graphics mode:	Display a "host is in graphics mode" message.
Bad text mode	Display a "host in an unsupported text mode" message.
Server off:	Display a "host is powered down" message.
Screen scroll up:	Scroll the screen up a certain number of lines.
Clear screen:	Clear the screen.
Move cursor:	Move the cursor.
Valid text mode:	Reflect valid text mode change.
Single character:	Display a character once.
Repeat character:	Display a character a certain number of times.
Display string:	Display a null terminated string.

The binary code format used for each packet type listed above is as follows:

Packet Type	Length	byte 0	byte 1	byte 2	byte 3	byte 4
Graphics mode:	1	0				
Bad text mode:	1	1				
Server off:	1	2				
Screen scroll up:	1	128+lines				
Clear screen:	2	0	attr.			
Move cursor:	2	row	column			
Valid text mode:	3	color/mono	rows	cols.		
Single character:	4	char	attr.	row	column	
Repeat character:	5	char	attr.	row	column	count
Display string:	6+	attr.	row	col.	char0	char1...

Bytes 3 and 4, and any subsequent bytes in the display string packet type contain a null terminated string to display.



This string must be at least two characters in length plus the null termination so that its length is greater than 5 and thus distinguishable from a repeat character packet using the length test. If the change only contains one character then the change is stored as a repeat character packet with a repeat count of one or a single character packet.

The function of the various buffers residing in the controller memory 32 shown in Fig. 5 will become more apparent in the discussion of the flowchart of Figs. 6a and 6b. Figs. 6a and 6b are flowcharts illustrating steps performed by processor 34 to process display data forwarded to buffers 48-52 as well as to store playback sequences amongst buffers 54-64.

Referring to Fig. 6a, step 70 illustrates processor 34 checking the local frame buffer 48 to see if any video mode changes have been made via the video control registers. Examples of video mode changes are a change between graphics and text mode or a change in the screen size, i.e., change in the number of rows and/or columns of a display unit. Before playback sequences can be stored, processor 34 must determine in step 72 if the current video mode is a valid text mode. If not, processor 34 waits until the video mode becomes a valid text mode before processor 34 copies the current display set (indicated by the start address specified in the video control registers) from the local frame buffer 48 to the temporary snapshot buffer 52 in step 74.

Processor 34 then compares the contents of the temporary snapshot buffer 52 with the contents of the snapshot buffer 50 to determine if a screen scroll has occurred in step 76. If a scroll has occurred, processor 34 stores a screen scroll packet to the appropriate buffers in step 78.

Processor 34 can then begin a sequence of steps to start at the beginning of the snapshot buffer 50 and temporary snapshot buffer 52 and compare them for changes. Processor 34 compares the current bytes of buffers 50 and 52 and looks for a change in step 80. If a change is found, as shown in step 82, then processor 34 updates in step 84 the snapshot buffer 50 with the change. Processor 34 then converts the change into a packet in step 86. It is the change, as applied to a packet, that is selectively placed into current buffers 54, 58 and 62 in step 88. Processor 34 then returns to step 80 to look for another change and repeats steps 82, 84, 86 and 88 until in step 82 processor 34 determines that there are no more changes between the display data in buffers 50 and 52.

Once processor 34 has determined that there are no more changes in step 82, processor 34 determines if the cursor has changed positions in step 90. If so, the processor 34 stores a move cursor packet into the current buffers 54, 58 and 62 in step 92 in a manner similar to that performed in step 88. If the cursor has not changed, then a packet rewrite to the existing cursor location is performed.

Referring now to Fig. 6b, a more detailed description of the steps taken to perform step 88 of Fig. 6a are shown. Processor 34 determines if the current reset sequence buffer 54 is full in step 94. That is, step 94 indicates if packets sent to buffer 54 has caused that buffer to become full. If buffer 54 is not full, then processor 34 forwards the present packet to the current reset sequence buffer 54 in step 96. Step 98 indicates whether processor 34 determines if current sequence buffer 58 is too full to receive another packet. If so, processor 34 removes the oldest packet from the current sequence buffer 58 in step 100, and thereafter converts the oldest packet into a character/attribute word. The character/attribute word of the removed packet is then forwarded to the current start buffer 62 in step 102. Processor 34 repeats steps 98, 100 and 102 until it determines in step 98 that the current sequence buffer 58 is not too full to receive the packet, at which time processor 34 puts the packet into the current sequence buffer 58 in step 104.

Following is a code fragment of C language function, find\_changes(), which can be used by processor 34 to periodically execute the functions set forth in the flowchart of Fig. 6a:

```

/*****
* find_changes()
5
* This routine is called repeatedly to find the screen changes and
* store them in the various buffers.
*****/
PRIVATE int find_changes(void)
10
{
    if (vga_mode_changed()) {

        /*-----
        VGA mode is different but may be in the middle of a change so delay
        for awhile to allow it to finish
15
        -----*/

        wait(DELAY_TIME);

        if (vga_mode_changed() & (VGA_UPDATE_MODE | VGA_UPDATE_ROWS |
20
            VGA_UPDATE_COLUMNS | VGA_UPDATE_VALID_STATUS)) {

            /*-----
            current VGA mode or screen dimensions have changed or the valid
            text mode status has changed so refresh screen
25

            NOTE: It is not necessary to store the entire screen if only the
            starting address changes.
            -----*/

            /* store the mode change and refresh the screen */
30
            refresh_screen();
        }
    }
    if (vga_valid_text_mode()) { /* if a valid text mode */
35
        /*-----
        copy the local frame buffer into the temporary snapshot buffer for
        comparison with the snapshot buffer

        NOTE: This copy is used for comparison rather than the actual
40
        local frame buffer to make fast scrolls more legible.
        -----*/

        memcpy(temp_snapshot_buf, local_frame_buffer, buffer_size);

45
        /*-----
        check if the screen has scrolled up
        -----*/

        scroll = detect_scroll();
50
        if (scroll) {
            /* store scroll */
            put_scroll_up(scroll, update_buffers);
55

```

```

    }

    /*-----
5      scan through the entire screen comparing the local frame buffer
      with the last snapshot taken, convert any differences to
      packets, and put these packets into the current destination buffers
      and update the snapshot buffer
    -----*/

10    vidasm_scan(temp_snapshot_buf,snapshot_buf,buffer_size,number_columns);

    /*-----
      update cursor position
    -----*/

15    /* read cursor position */
    error = vga_cursor_position(&cursor);

    /* compare with last updated position */
20    if (!error && ((cursor.row != cursor_row) ||
        (cursor.column != cursor_column))) {
        put_cursor(cursor.row,cursor.column,update_buffers);
    }
25    }
    return 0;
}

```

30 Following is a code fragment of C language function, put\_packet(), which can be used by processor 34 to periodically execute the functions set forth in the flowchart of Fig. 6b.

```

/*-----
35 * put_packet()
*
* This routine stores a packet in the reset and failure buffers.
*-----*/
PRIVATE void put_packet(char* packet,int length,int destination)
{
40    char temp_packet[MAX_PACKET_SIZE]; /* packet retrieved to make room */
    int temp_length; /* temp packet length */

    /*-----
45    store packet in buffers
    -----*/

    if (destination & RESET_SEQUENCE_BUFF) {

        /*-----
50    NOTE: this put does nothing if the buffer is full.
        -----*/

        buffer_put_packet(packet,length,reset_ptr);
    }
55    if (destination & FAILURE_SEQUENCE_BUFF) {
        while (buffer_put_packet(packet,length,failure_ptr)) {

```

```

5  /*-----
   buffer full so need to make room for the packet
   -----*/

   if (!buffer_get_packet(temp_packet,&temp_length,failure_ptr)) {
       /* put overflow in FAILURE_START_BUFF */
       play_packet(temp_packet,temp_length,&failure_start);
       failure_dirty_bits[current_failure] = 1;
10
       }
       else {
           /* failure buffer empty so can't make any more room */
           break;
15
       }
   }
}

```

20 Referring now to Fig. 7, a flowchart illustrating steps taken by processor 34 to replay a reset sequence of video screens on the remote terminal 28 is shown. The administrator first determines whether the current reset sequence or previous reset sequence is desired to be replayed. Processor 34 responds by choosing the current reset sequence buffer 54 or the previous reset sequence buffer 56, respectively, as the reset buffer to be replayed. When the system administrator instructs the server controller 26 to replay one of the reset sequences, the processor 34 determines in 25 step 106 if there are more packets in the reset buffer to display. If so, processor 34 retrieves the next packet from the reset sequence buffer and converts the packet to a character sequence in step 108. Processor 34 then displays the character sequence on the remote terminal 28 by transmitting the character sequence to the remote terminal 28 in step 110. Processor 34 repeats steps 106, 108, and 110 until processor 34 determines in step 106 that there are no 30 more packets to display.

Following is a code fragment of C language function, video\_reset\_play(), which can be used by processor 34 to periodically execute the functions set forth in the flowchart of Fig. 7.

```

35  /******
   * video_reset_play()
   *
   * This routine reads the next packet from a reset buffer and displays
40
45
50
55

```

```

*      it on the remote screen. Note that the packet is not removed from
*      the buffer. Call video_reset_rewind() to start playback from the
*      beginning of the buffer again.
5 *****/
PUBLIC int video_reset_play(void)
{
    char packet[MAX_PACKET_SIZE];
    int length;
10    int error = 0;
    unsigned long rc;

    /*-----
      read packet from buffer (does not remove the packet from the buffer)
15    -----*/

    if (buffer_peek_packet(reset_location, packet, &length, reset_buff)) {

        /*-----
          buffer empty so no packet received.
20        -----*/

        error = 0;
    }
25    else {

        /*-----
          buffer not empty so display the retrieved packet
          -----*/

30        reset_location++;          /* go to next location */
        play_packet(packet, length); /* display the packet */
        error = 1;
    }
35    return error;
}

```

Referring now to Fig. 8, a flowchart illustrating steps taken by processor 34 to replay a failure sequence on remote terminal 28 is shown. When the system administrator instructs the server controller 26 to replay the failure sequence, processor 34 determines in step 112 if there are more characters in the failure start buffer 64 to display. It is noted that current start buffer 62 and failure start buffer 64 contain characters (i.e., words comprising character/attribute pairs) instead of packets found within buffers (i.e., buffers 54-60) which store screen changes as opposed to the screens themselves. If there are more characters in buffer 64 to display, processor 34 retrieves the next character from the failure start buffer 64 and displays the character on remote terminal 28 by transmitting the character to the remote terminal 28 in step 114. Processor 34 repeats steps 112 and 114 until processor 34 determines in step 112 that there are no more characters in the failure start buffer 64 to display. It is noted that a character is not removed from the failure start buffer 64; instead, a copy of the character is taken.

Once processor 34 determines in step 112 that all of the characters in the failure start buffer 64 have been displayed, processor 34 searches for further packets in the failure sequence buffer 60 to be displayed in step 116. If there are no further packets, processor 34 has then concluded replaying the failure sequence. Otherwise, the processor 34 converts the packet to an ANSI character sequence in step 118. Processor 34 then displays the character sequence on remote terminal 28 by transmitting the character sequence to remote terminal 28 in step 120. Processor 34 repeats steps 116, 118 and 120 until it determines in step 116 that all of the packets in the failure sequence buffer 60 have been displayed.

Following is a code fragment of C language function, video\_failure\_play(), which can be used by processor 34 to periodically execute the functions set forth in the flowchart of Fig. 8:

```

/*****
***
5  *  video_failure_play()
   *
   *  This routine reads the next packet from the failure sequence buffer
   *  and
   *  displays it on the remote screen. Note that the first rows*columns
10  *  times this function is called, it displays the next character from
   *  the failure start buffer. After the entire failure start buffer has
   *  been
   *  displayed, packets are then read and displayed from the failure
   *  sequence
   *  buffer. Packets are not removed from the buffers.
15  *  Call video_failure_rewind() to restart the playback from the
   *  beginning.
   *****/
**/
20 PUBLIC int video_failure_play(display_screen* screen)
   {
       char packet[MAX_PACKET_SIZE];
       char character;
       int attribute,fore,back,rows,columns;
25       int length;
       int error;

30

35

40

45

50

55

```

```

        unsigned long rc;

        /*-----
5      -      set screen dimensions for 1st time through
        -----*/

        if ((failure_row == 1) && (failure_column == 1) && (failure_location ==
10      1)) {
            if (display_screen_dimensions(&last_failure_start,&rows,&columns) ==
0) {
                /* last_failure_start is valid */
                display_change_dimensions(rows,columns,screen);
15      }
        }

        /*-----
20      -      determine current display screen dimensions
        -----*/

        display_screen_dimensions(screen,&rows,&columns);
25      /*-----
        -      first display visible portion of the failure start buffer
        -----*/
30      /*

        while ((failure_row <= VGA_ROWS_MAX) &&
                (!display_row_visible(failure_row,screen))) {

35      /*-----
        -      while still within the buffer but the row is not visible

                Note: This loop is used to skip rows that are not visible, i.e.
                        the bottom (VGA_ROWS_MAX - 25) rows if the display screen
40      only      has 25 rows.
        -----*/

        failure_row++;
        failure_column = 1;
        }
        if (failure_buffer_dirty && (failure_row <= VGA_ROWS_MAX)) {

50      /*-----
        -      the failure start buffer still has changes which will be visible
on      the display screen. Read the next character from the buffer.
        -----*/
55      /*

```

```

character = VGA_READ_CHARACTER(failure_row, failure_column,
                                failure_start_buffer.params.vga);
attribute = VGA_READ_ATTRIBUTE(failure_row, failure_column,
                                failure_start_buffer.params.vga);
5   VGA_ATTRIBUTE_TO_COLOR(attribute, &fore, &back);

/*-----
10  -   display this character
-----*/

display_character(character, fore, back, failure_row, failure_column, 1,
15  screen);

/*-----
20  -   go to next column handling wrap-around
-----*/

failure_column++;
if (failure_column > columns) { /* end of line */
25  failure_column = 1;
    failure_row++;
}
error = 1; /* character displayed */
}
30  else {

/*-----
-   the failure start buffer has no more visible changes. Get next
35  packet from failure sequence buffer.
-----*/

if(buffer_peek_packet(failure_location, packet, &length, last_failure_ptr)){
40  /*-----
-   buffer empty so no packet received.
-----*/

45  /*
error = 0;
}
else {

50  /*-----
-   buffer not empty so packet received.
-----*/

55  /*
failure_location++;           /* go to next location */

```



```

        play_packet(packet,length,screen); /* display the packet */
        error = 1;
    }
5    }
    return error;
}

```

Playback of playback sequences begins when the administrator presses the space key in response to the menu shown in Fig. 2c. Playback consists of a stream of characters and escape sequences sent to the remote terminal to display the information in the sequence buffers. Each displayable character (ASCII values 32-126 and 128-254) is sent to remote terminal 28 as an 8-bit ASCII value. Some characters (ASCII values 0-31 and 127) are control characters that do such things as ring a bell rather than displaying an actual character. As a result, if one of these non-displayable characters appears in the local frame buffer 48, it is actually sent to the remote terminal as either a space or a similar looking displayable character. An escape sequence (a sequence of characters starting with the escape character (ASCII value 0x1B) is used to perform other actions such as moving the cursor, clearing the screen, and changing the current display colour. Server controller 26 transmits all characters unchanged except characters 0-31, 127, 255. These characters are remapped to the following ASCII bytes:

```

20      32, 32,32,32,32,32,32, 42, /* chars 0-7 */
      32, 32,32,32,32,32,32, 42, /* char 8-15 */
      62, 60,32,33,32,32,32, 32, /* chars 16-23 */
25      94,118,62,60,32,32,94,118, /* chars 24-31 */
      94, /* char 127 */
      32 /* char 255 */

```

The following are the escape sequences sent by the server controller 26 where \x1B = ESC character. In the event that the Compaq Insight Manager (CIM) application is executing on the remote terminal 28, server controller 26 transmits certain non-standard escape sequences to communicate with the CIM. Those non-standard escape sequences are denoted with an asterisk in the following Table I:

Table I

Always used:		n = cursor row, m = cursor column (i.e. \x1b[1;2H)
Clear screen:	\x1b[2J	
Move Cursor:	\x1b[n;mH	
*Flint ID:	\x1b{}	
-Color:		
+Disable hi intensity:	\x1b[0m	
+Foreground hi intensity:	\x1b[1m	
+Black Background:	\x1b[30m	
Red Background:	\x1b[31m	
Green Background:	\x1b[32m	
Yellow Background:	\x1b[33m	
Blue Background:	\x1b[34m	
Magenta Background:	\x1b[35m	
Cyan Background:	\x1b[36m	
+White Background:	\x1b[37m	
+Black Foreground:	\x1b[40m	
Red Foreground:	\x1b[41m	
Green Foreground:	\x1b[42m	
Yellow Foreground:	\x1b[43m	

Table I (continued)

Blue Foreground:	\x1b[44m	
Magenta Foreground:	\x1b[45m	
Cyan Foreground:	\x1b[46m	
+White Foreground:	\x1b[47m	
Combinations of above:	i.e. \x1b[0:1:47:30m	
Only used when talking to CIM:		
*Repeat Character:	c\x1b[nb	c = character to repeat. n = repeat count
*Screen Scroll:	\x1b[nS	n = number rows to scroll
*Set Number Rows:	\x1b[n~	n = number of rows
*Enable Japanese:	\x1b\$B\x1b(B	
*Disable Japanese:	\x1b[I	
-Color:		
*+Background hi intensity:	\x1b[5m	

The escape sequences marked with a '+' character in the above Table I are the only colour escape sequences used in monochrome modes. All colours except black and hi intensity black are converted to white or hi intensity white.

If the server controller 26 is not communicating with the CIM, the hi intensity background colours are not supported and are sent as the same colour but normal intensity. If the foreground is this same colour of normal intensity, then the foreground colour is remapped so that it will still be visible according to the following Table II:

Table II

Colour Modes	Monochrome Modes
Black -> Red	Black -> White
Red -> Black	White -> Black
Green -> Yellow	
Yellow -> Green	
Blue -> Magenta	
Magenta -> Blue	
Cyan -> White	
White -> Cyan	

Referring now to Figs. 9a-9e, screen shots of an exemplary failure sequence are shown. The screen sequence is exemplary of a playback sequence displayed by the server controller 26 on the remote terminal 28. Fig. 9a is a screen shot of a typical initial screen upon power-up of a server. The screen shows various items regarding the system configuration. Fig. 9b is a screen shot of a typical WINDOWS NT loader screen. The user is given the opportunity to select between a plurality of NT kernels to boot. Fig. 9c is a screen shot of a typical NT boot screen. The screen contains the build version of the kernel, the amount of system memory, processor configuration, etc. Fig. 9d illustrates the message displayed by the server controller 26 on the remote terminal 28 during a playback sequence when the playback sequence contains a packet indicating that the server 10 display 24 was placed in graphics mode by the host CPU 12. Figure 9e is a screen shot of a typical NT blue screen.

By snooping the expansion bus (preferably, EISA bus) for writes to the video subsystem, server controller 26 advantageously obtains video data updates in real time so that replay of the playback sequences can be realized. In addition, by snooping the expansion bus for writes to the video subsystem, server controller 26 obtains the video data updates in a manner typically non-intrusive to server 10 and, more specifically, non-intrusive to the expansion bus bandwidth. Other solutions obtain video data in an intrusive and non-real-time manner. For example, the COMPAQ Server Manager/R ("SMR") product asserts mastership of EISA bus cycles to obtain data within video memory 22 of server 10. SMR thereby periodically copies a portion of display data from video memory 22 into its own memory. The copied data is then compared with a previous copy of the video data to determine screen changes. The bus mastering approach of the SMR disadvantageously increases traffic on the EISA bus and adds extra bus arbitration delay to other EISA bus masters.

Referring again to Fig. 4, description of the various blocks 40-44 within detection logic 30 is now provided. Detection

logic 30 comprises address buffers 42, data transceivers 44, and control logic 40 coupled between the EISA bus and the controller bus. The control logic 40 and processor 34 are coupled by HOLD and HLDA (HOLD Acknowledge) signals. The address buffers 42, data transceivers 44, and control logic 40 cooperate to snoop display data and thereafter write the data to local frame buffer 48 (shown in Fig. 5) which forms a part of controller memory 32.

Controller memory 32 is a resource shared between processor 34 and detection logic 30, in that both the processor 34 and the detection logic 30 modify controller memory 32. The controller memory 32 comprises a single-ported pseudo-static RAM. Detection logic 30 therefore steals cycles away from processor 34 to write the video data to controller memory 32. As write cycles occur to the video subsystem, detection logic 30 decodes the write cycle, puts processor 34 on hold, and writes the video data into controller memory 32. Server controller 26 writes the video data to controller memory 32 concurrently with writing the video data to the video subsystem advantageously sustaining normal video activity on server 10.

Control logic 40 receives the address signals from the expansion bus and decodes the presence of an address in the video address ranges. If an address in the video address ranges appears on the expansion bus address signals, control logic 40 puts processor 34 on hold by asserting the HOLD signal to the processor 34 to arbitrate for controller memory 32. Processor 34 asserts HLDA to grant ownership of the controller bus to control logic 40.

Control logic 40 must ensure that it will be able to obtain ownership of the controller bus and write the video data to the controller memory 32 before completion of the write cycle to the video subsystem. To do so, control logic 40 makes use of the certain signals within an EISA expansion bus. The control logic utilizes EXRDY signal available to an EISA expansion bus. If the expansion bus is an ISA bus, then the control logic utilizes CHRDY signal. In either instance, EXRDY or CHRDY function to delay the expansion bus write cycle to the video subsystem until control logic 40 obtains ownership of the controller bus.

On the rising edge of the EISA expansion bus (expansion bus) BCLK signal, during assertion of the expansion bus START# signal, control logic 40 latches the address from the expansion bus address signals. Control logic 40 examines other control signals on the expansion bus, namely BCLK, START#, CMD#, and W\_R to determine if the bus cycle is a write cycle. If the bus cycle is not a write cycle, then control logic 40 stops arbitrating for the controller bus by deasserting the HOLD signal. If the bus cycle is a write cycle, data transceivers 44 drive the video data from the expansion bus data signals into the controller memory 32, thereby completing the write cycle on the rising edge of the expansion bus CMD# signal. The rising edge of CMD# is delayed if the control logic 40 drives the EXRDY signal low.

If the arbitration latency (the time between which control logic 40 asserts HOLD and processor 34 granting HLDA) exceeds the time needed to setup for a write cycle to the video address ranges, the control logic 40 drives the EXRDY signal low on the expansion bus to insert delay cycles, i.e., to delay the rising edge of CMD#. The EXRDY signal being driven low indicates to the master of the expansion bus which is currently performing the write cycle that the recipient of the write cycle is not ready to receive the data. Typically, processor 34 grants HLDA to control logic 40 sufficiently early so that the control logic 40 does not need to assert EXRDY. Thus, server controller 26 snooping is typically non-intrusive.

Control logic 40 translates the address previously latched from the expansion bus address signals to a corresponding location in the local frame buffer of controller memory 32, and also provide translated address to address buffers 42 which drive the translated address onto the controller bus. Control logic 40 then generates control signals to the data transceivers 44, address buffers 42 and controller bus to write the display data from the expansion bus data signals to controller memory 32 at the translated address driven by address buffers 42. Control logic 40 continues to keep processor 34 on hold until the end of a block transfer of display data, i.e., during burst or back-to-back write cycles to the video subsystem. In the event of a burst write cycle or back-to-back write cycles to the video subsystem, advantageously only the initial cycle has intrusion (i.e., incurs delay cycles), if at all.

Snooping write cycles to the video subsystem from the expansion bus into controller memory 32 requires synchronization between the expansion bus and the controller bus since they have different clock signals. This requires multiple synchronization points from both buses to guarantee correct states of signalling (i.e., to reduce the possibility of metastability). It is, therefore, advantageous to determine the earliest point at which a valid write cycle begins and the earliest point at which that write cycle, or series of write cycles, ends. It is therefore necessary to determine the earliest point at which a valid cycle is not a write cycle in order to negate HOLD, thus returning ownership of the controller bus back to the processor 34. This enables processor 34 to perform necessary functions such as processing the playback sequences. Additionally, it is advantageous to determine when a valid write cycle begins in order to, if necessary, drive EXRDY low and thereby insert delay cycles.

Expansion bus clock signal, BCLK, is not particularly useful in determining the beginning and ending of a write cycle because it is not guaranteed to be synchronous to either, i.e., the falling edge of START# or the rising edge of CMD#. For a detailed description of the timing specifications of an EISA bus, reference is made to EISA Specification Revision 3.12, which is hereby incorporated by reference.

The expansion bus Write/Read signal (W\_R), indicating a write cycle, is not valid as early as the expansion bus address signals (LA). According to EISA bus specifications, W\_R cannot be validated until the rising edge of BCLK as

START# is negating. Control logic 40 must be able to determine if a write cycle is occurring so that it does not respond to read cycles but instead allows the video subsystem to respond to read cycles.

EISA expansion bus signals START# and CMD# are not guaranteed to occur in sequence to each other as to the negation of START# and the assertion of CMD#. That is, the two signals may overlap or gap. Consequently, in order to use the falling edge of START# to denote the beginning of a cycle and the rising edge of CMD# to denote the end of a cycle, an equation wishing to denote the full cycle from start to finish must have some input to cover the potential gap between the negation of START# and the assertion of CMD#.

It was determined by the present inventors that in most bus interface implementations using, e.g., bus interface unit 18 as an EISA bus bridge, W\_R is valid as BCLK is low and START# is asserted low. This information is helpful in determining the beginning and end of a valid write cycle. In particular, the present invention employs a transparent latch control signal, N\_START, having the equation:

$$N\_START = ! BCLK * ! START\# + N\_START * CMD\#;$$

Thus, N\_START is asserted when BCLK is low and START# is asserted low, and stays asserted until CMD# goes asserted low to cover the previously discussed START# and CMD# gap.

An additional signal, EARLY\_DECODE, is defined which is true when the EISA expansion bus address signals are in the video address ranges.

Combining EARLY\_DECODE, W\_R, CMD# and N\_START a valid write cycle, is defined as:

$$VALID\_CYCLE = N\_START * EARLY\_DECODE * W\_R + VALID\_CYCLE * ! CMD\#;$$

Using VALID\_CYCLE, EXRDY is asserted upon VALID\_CYCLE being asserted and the processor 34 not yet having asserted HLDA. EXRDY is then held asserted until after the next falling edge of BCLK in which HLDA has been granted.

Thus, control logic 40 advantageously uses N\_START to meet the EISA timing requirements with respect to EXRDY and thus in typical situations to cause no intrusion, i.e., no introduction of delay cycles, on operations to the video subsystem. Control logic 40 also uses N\_START to perform synchronization between the EISA expansion bus and the controller bus since the two buses have different clock signals. By using N\_START (i.e., the term with BCLK being low), the control logic 40 advantageously enjoys at least a 60 nanosecond advantage in synchronization time over using the rising edge of START# (which occurs at least 60 nanoseconds later, i.e., synchronous with the rising edge of BCLK) to determine a valid write cycle.

Referring now to Fig. 10, a timing diagram illustrates the relationship of the EISA expansion bus signals BCLK, START#, CMD#, LA, W\_R to the EARLY\_DECODE, HOLD, N\_START, VALID\_CYCLE, and EXRDY signals. Fig. 10 depicts a video write cycle in which processor 34 does not grant HLDA sufficiently early enough and the control logic 40 must drive EXRDY low to inset delay cycles followed by a second video write cycle. The first and second video write cycles constitute back-to-back video write cycles. According to the present invention, the second video write cycle advantageously does not incur a delay cycle even though the first write cycle did incur a delay cycle. The edges of BCLK are numbered consecutively one through fourteen for clarity in the discussion of Fig. 10. HLDA signal shown in Fig. 10 is a version of the HLDA signal from processor 34 which has been synchronized to the EISA expansion bus clock BCLK.

EISA bus interface 18, on behalf of CPU 12 wishing to write to the video subsystem, generates an address in the video address ranges on the LA signals near BCLK falling edge two. The control logic 40 responsively asserts the EARLY\_DECODE signal. Shortly thereafter, the control logic 40 asserts HOLD to the processor 34. The HOLD signal is asserted when the EARLY\_DECODE signal becomes asserted and remains asserted as long as the VALID\_CYCLE signal remains asserted. By keeping HOLD asserted through VALID\_CYCLE, the control logic 40 keeps ownership of the controller bus until the write of the video data to the controller memory 32 completes. Additionally, the HOLD signal is asserted as long as EARLY\_DECODE is asserted and the defeat timer has not expired (as discussed below). The HOLD signal is synchronized with processor 34 clock. Near BCLK rising edge three, bus interface 18 asserts START# (which is active low) to begin the address phase of the cycle. The bus interface 18 asserts the W\_R signal to signify a write cycle.

In response to BCLK falling edge four, the control logic 40 asserts N\_START since BCLK is low and START# is asserted low. N\_START remains asserted until the assertion of CMD# low near BCLK rising edge five. N\_START being active signifies that W\_R is valid and since W\_R indicates a write cycle and EARLY\_DECODE indicates an address in the video address ranges, the control logic 40 correspondingly asserts VALID\_CYCLE. In response to VALID\_CYCLE being asserted after BCLK falling edge four, the control logic 40 drives EXRDY low since HLDA has not been granted (in this example). In this example, HLDA is granted near BCLK rising edge five. The control logic 40 continues to drive

EXRDY low until after the next falling edge of BCLK in which HLDA is granted, i.e. BCLK falling edge six since EXRDY is sampled by bus interface unit 18 on the falling edge of BCLK while CMD# is asserted low.

Bus interface unit 18 asserts CMD# low and deasserts START# high after BCLK rising edge five to begin the data phase of the write cycle. As previously noted, the rising edge of START# and the falling edge of CMD# here may either gap or overlap according to the EISA expansion bus specification. As previously described, N\_START advantageously covers the gap so that VALID\_CYCLE may denote a valid video write cycle from start to finish.

In response to EXRDY being driven low during the falling edge of BCLK during CMD# asserted low, i.e., BCLK falling edge six, bus interface unit 18 waits an extra BCLK cycle, i.e., inserts a delay cycle (or wait state), by delaying the deassertion of CMD# high for an extra clock cycle near BCLK rising edge nine. CMD# deasserting high denotes the end of the write cycle.

Since the address on the LA signals remains in the video range for the second write cycle, control logic 40 continues to assert EARLY\_DECODE and correspondingly to assert HOLD thus retaining ownership of the controller bus as signified by HLDA being asserted through the end of the second write cycle.

After BCLK rising edge nine, bus interface unit 18 (using EISA bridging) asserts START# to signify the beginning of the address phase of the second video write cycle. The second write cycle is similar to the first video write cycle except that since the control logic already has ownership of the controller bus EXRDY is not negated low during the second write cycle. Therefore, bus interface unit 18 deasserts CMD# high near BCLK rising edge thirteen (i.e., a full clock cycle earlier than in the first write cycle) to end the data phase of the second write cycle.

The second video write cycle is also representative of a non-back-to-back video write cycle in which HLDA was granted sufficiently early such that control logic 40 did not have to drive EXRDY low to insert delay cycles.

In one embodiment of the present invention, server controller 26 further comprises a controller video subsystem (not shown) similar to the video subsystem of server 10 comprising video controller 20 and video memory 22. The controller video subsystem may be selectively disabled by system configuration software. The controller video subsystem is an ISA device. If the controller video subsystem is enabled, server controller 26 negates the ISA bus signal CHRDY rather than the EISA expansion bus EXRDY signal.

A purpose of the controller video subsystem is to cause writes of video data to appear on the expansion bus in the event that the VGA controller on server 10 is a PCI VGA controller. Otherwise, the PCI VGA controller would decode and respond to video writes thus causing the bus interface unit 18 to not forward the video write to the EISA expansion bus. As a result, there would be no video write cycles on the EISA expansion bus or the server controller 26 to snoop. Thus, when a PCI VGA controller resides in the system, the user must enable the controller video subsystem on the server controller 26 and disable the PCI VGA controller in order to enable video write cycle snooping and playback sequence capturing.

It is common for implementations of bus interface unit 18 (i.e., EISA bus bridge) to drive an address on the expansion bus address signals to, at the end of a valid cycle, latch and drive that address on the expansion bus (i.e., EISA bus) during an idle cycle occurring after the valid cycle, so the address signals do not float. This condition would cause control logic 40 to assert the EARLY\_DECODE signal and consequently to assert HOLD to the processor 34 continuously for relatively large amounts of time. As a result processor 34 would be needlessly starved. One solution to this problem would be to negate HOLD at the end of a cycle on the rising edge of CMD#. However, the result would be that the control logic 40 would have to re-arbitrate for the controller bus and re-synchronize between the EISA expansion bus and controller bus for each write cycle of a block of back-to-back write cycles by negating EXRDY thus introducing delays.

The present invention employs a defeat timer in order to solve the problem described above. Control logic 40 comprises a counter which counts a certain number of BCLKs after a cycle if the EISA expansion bus address signals still have an address in the video address ranges. When the counter reaches terminal count without a new assertion of START#, the control logic 40 negates HOLD and does not assert HOLD until the next START# assertion with the control logic 40 decoding an address in the video address ranges.

Thus, by employing the defeat timer, control logic 40 advantageously asserts HOLD earlier than would otherwise be possible if the control logic 40 waited until the assertion of START# to assert HOLD and is thus more likely to not introduce delay cycles to the server 10 during video data writes to the video subsystem. The defeat timer further advantageously enables control logic 40 to keep HOLD asserted during a block of back-to-back valid video write cycles, thus potentially avoiding introducing delay cycles to server 10 during video data writes to the video subsystem.

Remote terminal 28 establishes a communications connection with server controller 26 via communication unit 36. The connection between the remote terminal 28 and server controller 26 is commonly referred to as an out-of-band connection, or asynchronous connection. An out-of-band connection refers to a network connection that is established through a phone line or direct serial connection, rather than through a standard network medium such as a local area network Ethernet connection. By providing an out-of-band connection, server controller 26 can establish communication with the remote terminal 28 both when server 10 is functioning normally or when a failure occurs, such as an operating system crash or a failure on the local area network to which server 10 is connected.

In a first mode, remote terminal 28 establishes connection with the server controller 26 using the ANSI terminal emulation protocol, referred to henceforth as the text protocol. In a second mode, remote terminal 28 establishes a point-to-point-protocol (PPP) connection with server controller 26. PPP is described in numerous Request For Comments (RFC) of the Internet Engineering Task Force ("IETF") whose Uniform Resource Locator ("URL") is <http://www.ietf.cnri.reston.va.us/home.html>. In particular RFC 1661, 1662 and 1663 are pertinent to the present invention and are hereby incorporated by reference. When the remote terminal 28 dials in to the server controller 26, server controller 26 software executing on the processor 34 of the server controller 26 advantageously automatically determines which of the two data link layer protocols the remote terminal 28 is using.

In the first mode, a portion of server controller 26 software, referred to as the server controller console application, communicates with remote terminal 28 using the text protocol. The server controller console application enables remote terminal 28 to access various features of the server controller 26 such as viewing playback sequences. The server controller console application additionally allows the system administrator to perform remote resets, manage alert and login information, remotely access the console, view event logs, error logs and status logs of server 10. Referring again to Figs. 2a-2c, three menu screens of the server controller console application are shown. Typically, remote terminal 28 is a computer executing a popular ANSI terminal emulation software application available from various vendors, such as the well known products SYMANTEC pcANYWHERE or DATASTORM TECHNOLOGY ProComm, to communicate with server controller 26.

In the second mode, remote terminal 28 establishes a TCP/IP link with server controller 26 over the PPP connection similar to that of a local area network (LAN). Typically, remote terminal 28 executes the COMPAQ Insight Manager (CIM) to communicate with the server controller 26 over the TCP/IP connection. Server controller 26 communicates with the CIM in two manners. First, server controller 26 communicates with the CIM via a TELNET connection to execute the server controller console application just as in the first mode using the text protocol.

Second, server controller 26 sends/receives SNMP packets to/from the CIM for the purposes of managing the server 10. For example, server 10 may detect a failure or other event on server 10 and send an SNMP trap packet through server controller 26 to the CIM executing on the remote terminal 28 to notify the system administrator of the event just as it would if the CIM were executing on a client on the LAN to which server 10 is connected. Additionally, in the event of a server 10 failure, server controller 26 may detect the failure and autonomously send an SNMP trap packet to the CIM. In this second mode, i.e., over the PPP connection, server controller 26 may simultaneously communicate with the remote terminal 28 via the server controller console application and SNMP packets.

Another example of server 10 management via CIM is the remote terminal 28 sending a "get request" SNMP packet to the server 10 via server controller 26 to request management information about the server, such as the health of one of the server's subsystems. Server controller 26 forwards the SNMP packet to the operating system. The operating system gathers the requested information and sends a "get response" packet to the server controller 26. The server controller 26 in turn forwards the packet to the CIM.

The WINDOWS NT Remote Access Services (RAS) can establish PPP connections to provide communication between two computers connected by a null-modem. The NT RAS employs a proprietary method of establishing this PPP connection. The method, referred to herein as the "pre-PPP protocol", is to exchange a pair of strings between the two computers. The first computer, wishing to establish the connection, sends the string "CLIENT" to the second computer. Next, the second computer, after receiving the "CLIENT" string, sends the string "CLIENTSERVER" to the first computer. Once this exchange of the two strings has occurred, the two computers proceed to exchange the necessary character packets to establish the PPP connection. In standard PPP connections this initial exchange of the two strings is not required. Server controller 26 is capable of establishing a PPP connection with both the NT RAS over a null-modem as well as remote terminal 28 using the standard PPP protocol.

Referring now to Fig. 11, a flowchart is presented which illustrates how server controller 26 detects which communication protocol is sent from remote terminal 28. Prior to remote terminal 28 establishing a connection, server controller 26 software resets three state machines within the server controller 26 software in step 130. A state machine exists for each of the valid protocols: the text protocol, the PPP protocol, and the pre-PPP protocol.

During times when server controller 26 waits for an interrupt from communication unit 36, playback sequences can be processed over by processor 34. When a carrier signal is received upon server controller 26, software within controller 26 resets the state machines in step 134 and sets a timer in step 136. In one embodiment, the timer is set to expire in 30 seconds. The server controller 26 software determines if the timer has expired in step 138. If the timer has expired, the server controller 26 software declares the protocol to be the text protocol by communicating to an upper level task in the server controller 26 software that the protocol is the text protocol in step 140.

Prior to receiving a character or signal from remote terminal 28, software within server controller 26 can process the video screen changes. Once the interrupt occurs indicative of a character or signal present, server controller 26 software determines if a "no carrier" signal was sent by the communication unit 36 to indicate a loss of carrier from the remote terminal 28 in step 142. If the carrier has been lost, the server controller 26 software clears the timer in step 144 and returns to step 130.

If a "no carrier" signal was not received in step 142, the server controller 26 software determines if a character was received in step 146. If not, the server controller 26 software returns to step 138 to wait for either the timer to expire or an interrupt to be received indicating that a signal or character has arrived.

If a character was received in step 146, the server controller 26 software passes the character to each of the three protocol state machines in step 148. If none of the state machines detect a valid protocol in step 150 the server controller 26 software returns to step 138. If one of the state machines detects a valid protocol in step 150, the server controller 26 software clears the timer in step 152 and then declares to the upper level task which of the three protocols was detected in step 154.

The processor 34 is configured to receive interrupts from various sources, such as the communication unit 36 as mentioned in the description of Fig. 10 above. The processor 34 comprises an interrupt descriptor table, wherein each entry in the table references an interrupt service routine associated with an interrupt source. When one of the interrupt sources generates an interrupt, processor 34 begins executing the interrupt service routine referenced by the entry in the interrupt descriptor table corresponding to the interrupt source. The interrupt service routines are part of the server controller 26 software and execute on the processor 34. Communication unit 36 is configured to interrupt processor 34 upon reception of a character or signal, such as a "carrier" or "no carrier", from the remote terminal 28. The interrupt service routines communicate with other software tasks executing on the processor 34 to facilitate communications with the remote terminal 28.

Initially, i.e., prior to the remote terminal 28 establishing a connection with the server controller 26, the interrupt descriptor table entry associated with the communication unit 36 is populated with a reference to a protocol detection interrupt service routine. The protocol detection interrupt service routine is the portion of the server controller 26 software which comprises the three protocol state machines. The protocol detection interrupt service routine is also the portion of server controller 26 software which receives an incoming character from communication unit 36. Each character received by the protocol detection interrupt service routine acts as the input to the three protocol state machines.

Once the upper level task has been notified that a valid protocol has been detected, either in step 140 or step 154, the upper level task populates the interrupt descriptor table entry associated with the communication unit 36 with a reference to an interrupt service routine specific to the particular protocol detected. The PPP and pre-PPP protocols have the same interrupt service routine.

Appendix A (see below) contains an assembly language source code listing of portions of the PPP interrupt service routine. The code portions illustrate how the interrupt service routine receives and processes incoming characters to validate and/or form a valid and whole PPP packet. Once a PPP packet has been formed, the interrupt service routine notifies an upper level task of server controller 26 software regarding the presence of that packet.

Referring now to Fig. 12, a state diagram illustrating the text protocol state machine is shown. When the text protocol state machine detects three consecutive ASCII "return" characters (0x20), it indicates that a valid text protocol has been detected from the remote terminal 28. In the TEXT\_INIT state 156, if the state machine receives a carriage return character, the state machine transitions to the TEXT\_ONE state 157. Otherwise the state machine remains in the TEXT\_INIT state 156. In the TEXT\_ONE state 157, if the state machine receives a carriage return character, the state machine transitions to the TEXT\_TWO state 158. Otherwise the state machine returns to the TEXT\_INIT state 156. In the TEXT\_TWO state 158, if the state machine receives a carriage return character, the state machine transitions to the TEXT\_DETECTED state 159. Otherwise the state machine returns to the TEXT\_INIT state 156. When the state machine reaches the TEXT\_DETECTED state 159, it notifies an upper level task in the server controller 26 software that a valid text protocol has been detected.

Referring now to Fig. 13, a state diagram illustrating the pre-PPP state machine is shown. When the pre-PPP state machine detects the string "CLIENT", it indicates that a valid pre-PPP protocol has been detected from the remote terminal 28 and correspondingly transmits the string "CLIENTSERVER" to the remote terminal 28. In the CLIENT\_INIT state 162, if the state machine receives a 'C', the state machine transitions to the CLIENT\_C state 163. Otherwise the state machine remains in the CLIENT\_INIT state 162. In the CLIENT\_C state 163, if the state machine receives a 'L', the state machine transitions to the CLIENT\_L state 164. This progression is maintained for each of the subsequent states 164, 165 and 166 to indicate possible receipt of CLIENT string. If the string is not fully received, then a transition occurs back to CLIENT\_INIT state 162. Upon receiving the entire string, the state machine transitions to CLIENT\_DETECTED state 168 and notifies an upper level task in the server controller 26 software that a valid pre-PPP protocol has been detected.

The following illustrates the format of the character sequence forming a valid PPP packet:

SYNCH ADDRESS CONTROL PROTOCOL ... data ... CRC<sub>low</sub> CRC<sub>high</sub> SYNCH

The SYNCH character is designated by a hexadecimal value 0x7E, the ADDRESS character is designated by a hexadecimal value 0xFF, and the CONTROL character is designated by a hexadecimal value 0x03. The ADDRESS and CONTROL character are optional, but if one is present both characters must be present. The PROTOCOL character (s) is one or two bytes depending upon the least significant bit (LSB) of the first byte. If the first byte has LSB set (i.e., 1), then the PROTOCOL character consists of only one compressed protocol byte. If the first byte has the LSB clear

(i.e., 0), then the PROTOCOL characters consist of two uncompressed protocol bytes, the first byte being the low byte and the second byte being the high byte. The "data" bytes are the payload of the packet. The CRC<sub>low</sub> and CRC<sub>high</sub> characters comprise low and high bytes, respectively, of a cyclic redundancy check (CRC) on the data bytes.

Referring now to Fig. 14, a state diagram illustrating the PPP state machine is shown. When the PPP state machine detects a valid PPP packet, it indicates that a valid PPP protocol has been detected from the remote terminal 28. In the PPP\_INIT state 170, if the state machine receives a SYNCH character, the state machine transitions to the PPP\_SYNCH state 172. Otherwise the state machine remains in the PPP\_INIT state 170. In the PPP\_SYNCH state 172, if the state machine receives an ADDRESS character, the state machine transitions to the PPP\_ADDRESS state 174. Otherwise the state machine transitions to the PPP\_CONTROL state 176. In the PPP\_ADDRESS state 174, if the state machine receives a CONTROL character, the state machine transitions to the PPP\_CONTROL state 176. Or, if the state machine receives a SYNCH character, the state machine transitions to the PPP\_SYNCH state 172. Otherwise the state machine transitions to the PPP\_INIT state 170. In the PPP\_CONTROL state 176, if the state machine receives a SYNCH character, the state machine transitions to the PPP\_SYNCH state 172. Otherwise, if the LSB of the character is clear the state machine transitions to the PPP\_NOT\_COMPRESSED state 178 treating the character as the low byte of two non-compressed protocol bytes and if the LSB of the character is set the state machine transitions to the PPP\_COMPRESSED state 180 treating the character as the single compressed protocol byte. In the PPP\_NOT\_COMPRESSED state 178, if the state machine receives a SYNC character, the state machine transitions to the PPP\_SYNCH state 172. Otherwise, the state machine transitions to the PPP\_COMPRESSED state 180 treating the character as the high byte of two non-compressed protocol bytes. In the PPP\_COMPRESSED state 180, if the state machine receives a SYNCH character, the state machine transitions to the PPP\_DETECTED state 182. Otherwise, the state machine remains in the PPP\_COMPRESSED state 180 and treats the character as a data or CRC byte. When the state machine reaches the PPP\_DETECTED state 182, it notifies an upper level task in the server controller 26 software that a valid PPP protocol has been detected.

Server controller 26 provides an optional dial-back security feature to prevent unauthorized users from tampering with the server 10 via server controller 26. Using the server controller 26 console application software, the system administrator configures a list of user profiles each containing a username, password and phone number. When a user dials in, the server controller 26 receives a username and password from the user, verifies the username and password from the list of user profiles and then disconnects. If the username and password are verified in the list of user profiles, the server controller 26 then dials up the associated phone number in the list of user profiles. For an authorized user, the remote access proceeds as usual. If an intruder had obtained a valid username and password, the intruder will nevertheless lose the connection to the server 10 since the intruder will not be dialled back.

Server controller 26 advantageously dials back the authorized user with the same communications protocol used when the user dialled in. For example, if the user connected via a null-modem from a WINDOWS NT machine using the pre-PPP protocol, the server controller 26 will dial back the NT machine using the pre-PPP protocol. This is advantageous since it is likely that the remote terminal will expect to communicate with the server controller 26 using the same protocol with which it dialled in.

Referring now to Fig. 15, various components of server controller 26 are shown arranged upon a PCB 200. PCB 200 is segregated into a plurality of planes or portions. A first portion 202 accommodates an expansion bus interface unit 204 as well as detection logic 30. A second portion 206 accommodates processor 30 and controller memory 32, as well as a secondary power supply (i.e., battery) denoted as reference numeral 208. A third portion 210 accommodates communication unit (i.e., modem) 36. Portions 202, 206 and 210 are segregated from each other by a pair of decoupling units denoted as first decoupling unit 212 and second decoupling unit 214.

Power to PCB 200 is normally applied via edge connectors 216 from the expansion bus to which PCB 200 can be connected. One pin of edge connector 216 is reserved for primary power VDD routed to first, second and third portions 202, 206 and 210, respectively. Display data is snooped from the expansion bus and is forwarded over possibly several pins of edge connector 216 to detection logic 30 via interface unit 204. Upon losing VDD during, for example, server power down, decoupling unit 212 serves two functions. First, power conductors within first portion 202 are decoupled from power conductors within second portion 206. More specifically, the VDD conductor within first portion 202 is decoupled via decoupling unit 212 from the power conductor or conductors within second portion 206. Second, signal conductors within first portion 202 are decoupled from signal conductors within second portion 206. As the signal conductors are decoupled, the signal conductors are tri-stated.

Similar to the first decoupling unit 212, second decoupling unit 214 also decouples the power and signal conductors. If VDD ceases from the expansion bus and if communication unit 36 does not detect communication from remote terminal 28, then second decoupling unit 214 performs its decoupling operation. If, however, VDD is present or if remote terminal 28 is sending communication protocol, then second decoupling unit 214 effectuates coupling between second and third portions 208 and 210, respectively.

Decoupling units 212 and 214 preferably comprise quick switches upon the power conductors and bi-directional transceiver with buffer/driver capability between signal conductors. The quick switches are obtainable from various



vendors, such as National Semiconductor Corp. and Texas Instruments, Inc., and are formed as packaged arrays of JFETs with a common control input (i.e., output-enable OE input). The bi-directional transceiver are well suited for driving high-capacitive loads and operate under high-speed conditions. A suitable bi-directional transceiver can be obtained from Integrated Device Technology, Inc. or Texas Instruments, Inc. as fast CMOS transceiver with tri-state output capability.

Coupled to one pin of edge connector 216 is a charge conductor 220. Charge conductor 220 is adapted to receive a voltage supplied by the expansion bus defined to exceed VDD. According to one embodiment, the desired voltage upon charge conductor 220 is 12 volts, wherein VDD equals 5 volts. Voltage supplied by charge conductor 220 serves not only to charge battery 208, but also to enable switching of switching regulator 222. If, for example, charge within charge conductor 220 goes below a threshold value, then switching regulator 222 activates coupling of battery 208 to second portion 206, and possibly third portion 210. Conversely, if charge on charge conductor 220 exceeds the threshold value, then output from battery 208 (i.e., VDDDB) is not coupled to second portion 206 and third portion 210. Instead, in the latter instance, VDD is present thereto. A diode 224 is used to prevent discharge of battery 208 during times when charge upon conductor 220 is terminated.

A ring indicator 226 is coupled on the same power plane as second portion 206. Specifically, ring indicator receives VDDDB upon conductor 228 during times when VDD is not present, or receives VDD if VDD is present. Accordingly, ring indicator 226 maintains power either from the primary or the secondary sources to detect an incoming analog ring ARNG signal. Upon receiving ARNG, ring indicator 226 forwards the ensuing communication protocol to communication unit 36 via the digital ring DRNG signal. By always maintaining power to ring indicator 226, detection of an incoming ring is always possible regardless of VDD status.

Fig. 16 illustrates a comparator 230 used in determining if VDD becomes less than VDDDB by a threshold amount. If VDD should happen to drop beyond the acceptable threshold, then comparator 230 produces a VDDOFF signal. The VDDOFF signal is presented on the output enable OE pin of decoupling of first decoupling unit 212. Comparator 230 thereby effectuates decoupling of signal conductors and power conductors between the first and second portions 202 and 206 if VDD supplied by the expansion bus becomes less than VDDDB supplied by battery 208.

Fig. 17 illustrates operation of digital ring DRNG, and its effect upon decoupling of power and signal conductors between second and third portions 206 and 210, respectively. If power to communication unit 36 is initially off, DRNG signal to communication unit 36 causes an interrupt to processor 34. As such, processor 34 is informed that a ring signal is present at communication unit 36, whereby processor 34 can then acknowledge its interrupt through MDMACT signal as well as presenting power to the modem via MDMPWR signal. More specifically, MDMPWR signal is fed to the output enable OE input of second decoupling unit 214. Upon receiving MDMPWR indicative of an incoming ring, second decoupling unit 214 couples the power and signal conductors extending between second portion 206 and third portion 210. Essentially, DRNG is gated with a signal indicative of whether communication unit 36 is on or off, the output of that gate being an interrupt to processor 34. Whenever processor 34 receives an interrupt, power from VDDDB of processor 34 is supplied to communication unit 36. If power to communication unit 36 is on, then an interrupt is not asserted.

It will be appreciated to those skilled in the art having the benefit of this disclosure that this invention is capable of applications with numerous types of computer systems, with numerous types of electrical components, with numerous types of PCBs and with numerous types of communications protocols. Furthermore, it is to be understood that the form of the invention shown and described is to be taken as presently preferred embodiments. Various modifications and changes may be made without departing from the spirit and scope of the invention as set forth in the claims. It is intended that the following claims be interpreted to embrace all such modifications and changes and, accordingly, the specification and drawings are to be regarded in an illustrative rather than a restrictive sense.

APPENDIX A

Unpublished Work  
 Copyright © 1996  
 Compaq Computer Corp.  
 All Rights Reserved

```

; States for receive state machine
10 RS_STATE_1 equ 1 ; waiting for first SYNCH
    RS_STATE_2 equ 2 ; waiting for another SYNCH or ADDR
    RS_STATE_3 equ 3 ; got ADDR; waiting for CONTROL
    RS_STATE_4 equ 4 ; read first PROTOCOL byte
    RS_STATE_5 equ 5 ; read second PROTOCOL byte
15 RS_STATE_6 equ 6 ; read data, adjusting CRC and waiting for
    SYNCH
    RS_STATE_7 equ 7 ; waiting to get a packet buffer
    RS_STATE_ESC equ 80h

; For now, hardcoded modem registers for COM3 (16550)
20 THR equ 03e8h ; Transmitter Holding Register
    RBR equ 03e9h ; Receiver Buffer Register
    IER equ 03e9h ; Interrupt Enable Register
    IIR equ 03eah ; Interrupt Identification Register
    LCR equ 03ebh ; Line Control Register
25 MCR equ 03ech ; Modem Control Register
    LSR equ 03edh ; Line Status Register
    MSR equ 03eeh ; Modem Status Register

; These are some bit positions/masks of some of the registers above.
; They are used to assert/clear certain signals.
30 IER_XMIT equ 02h ; Transmit interrupt enable bit in IER
    LSR_CHAR_AVAIL equ 01h ; Char available bit in LSR
    LSR_XMIT_EMPTY equ 20h ; THR empty bit in LSR
    MSR_DELTA_CD equ 08h ; Delta CD bit in MSR
    MSR_CD equ 80h ; CD bit in MSR
35 MCR_RTS equ 02h ; RTS bit in MCR
    MSR_CTS equ 10h ; CTS bit in MSR

; These are in other files, but are not necessary to
; understand this file.
40 extrn crc16_table:word ; 16-bit CRC lookup table
    extrn ppp_buffer_get_rcv:near ; get a packet buffer
    extrn ppp_buffer_ret_rcv:near ; return a packet buffer
    extrn ppp_q_write:near ; inform upper layer of packet reception

rcv_done:
45 call ppp_com3_read ; read incoming char into AL

    cmp al,PPP_SYNCH_CHAR ; is it a SYNCH char ?
    jz process_synch ; yes, handle these offline

    test ah,RS_STATE_ESC ; was last char an escape ?
50 jnz escaped ; yes, process offline

    cmp al,PPP_ESCAPE_CHAR ; is current char ESC?
    jz do_escape ; yes, take care of it

process_char:
55 xor bl,al ; we have received a true char
    movzx esi,bl ; modify the running
    ; CRC value to reflect

```

```

    shl     bx,8                ; the char that just
    xor     bx,crc16_table[esi*2] ; came in

    cmp     ah,RS_STATE_6       ; read a char of data
    jz      rx_state_6         ; (most common state)

    cmp     ah,RS_STATE_5       ; getting 2nd half of protocol
    jz      rx_state_5

    cmp     ah,RS_STATE_4       ; getting valid protocol
    jz      rx_state_4

    cmp     ah,RS_STATE_3       ; waiting for CONTROL
    jz      rx_state_3

    cmp     ah,RS_STATE_2       ; waiting for ADDRESS, CONTROL
    jz      rx_state_2         ; or valid protocol

    cmp     ah,RS_STATE_1       ; waiting for SYNCH, so
    jz      rcvr_done          ; throw away this char

    cmp     ah,RS_STATE_7       ; waiting for a buffer
    jz      rx_state_7

    mov     ah,RS_STATE_1       ; if we fell through there's a
                                ; problem...
    jmp     rcvr_done

rx_state_2:
    mov     edi,ppp_com3_rx_buffer_start ; re-initialize these two
                                ; variables
    mov     ebx,PPP_FCS16_INITIAL_VALUE ; so we always start in
                                ; a known state
    xor     bl,al              ; modify the running
    movzx   esi,bl             ; CRC value to reflect
    shr     bx,8               ; the char that just
    xor     bx,crc16_table[esi*2] ; came in

    cmp     al,PPP_ADDRESS_FIELD ; is it the ADDRESS field ?
    mov     ah,RS_STATE_4       ; assume no
    jnz     rx_state_4         ; good guess
    mov     ah,RS_STATE_3       ; cops, bad guess
    jmp     rcvr_done          ; go get CONTROL field

rx_state_3:
    cmp     al,PPP_CONTROL_FIELD ; did we get CONTROL ?
    mov     ah,RS_STATE_4       ; assume yes
    jz      rcvr_done          ; good guess
    mov     ah,RS_STATE_1       ; wrong. since there is no
    jmp     rcvr_done          ; protocol FF, reset state
                                ; machine

rx_state_4:
    movzx   ax,al              ; expand into AX
    test    al,1               ; compressed protocol ?
    jz      high_part          ; no
    mov     word ptr ppp_com3_rx_protocol,ax ; yes, so store it
    mov     ah,RS_STATE_6       ; and go to
    jmp     rcvr_done          ; the next state

high_part:
    shl     ax,8               ; we only got the high 8 bits
    mov     word ptr ppp_com3_rx_protocol,ax ; store the
                                ; temporary result
    mov     ah,RS_STATE_5       ; and go to
    jmp     rcvr_done          ; the next state

```

```

rx_state_1:
    test    al,1                ; this had better be set....
    jz      bad_protocol        ; bad packet header
    mov     byte ptr ppp_com3,,rx_protocol,al ; everything's
5                                     ; cool,
                                     ; save it
    mov     ah,RS_STATE_6       ; and go to
    jmp     rcvr_done           ; the next state
bad_protocol:
    mov     ah,RS_STATE_1       ; reset state machine
10    jmp     rcvr_done         ; and get another char

rx_state_6:
    stosb                       ; we have a valid data char
                                ; so store it
                                ; and fall through (it's a
                                ; given)
15    jmp     rcvr_done

rx_state_7:
    call    ppp_buffer_get_rcvr ; see if we can get a buffer
    mov     edi,eax             ; save result in EDI
    inc     eax                 ; turns -1 into 0
20    or     eax,eax            ; did we get a buffer ?
    jz      rcvr_done           ; no, stay in this state
    mov     ppp_com3_rx_buffer_start,edi ; yes, reset the pointer
    mov     ah,RS_STATE_1       ; and go to state 1
    jmp     rcvr_done           ; waiting for a synch

25    escaped:
        and     ah,07fh         ; clear RS_STATE_ESC flag,
        xor     al,PPP_ESCAPE_XOR_MASK ; xform this char,
        jmp     process_char    ; then process it normally

do_escape:
30    or     ah,RS_STATE_ESC    ; upon receiving an ESC
                                ; we set the RS_STATE_ESC
                                ; flag
        jmp     rcvr_done       ; and drop this char

process_synch:
35    cmp     ah,RS_STATE_6     ; we are inside of a packet
    jz      synch_6            ; this state is the most
                                ; common one encountered
                                ; under normal conditions
    cmp     ah,RS_STATE_2     ; we already got a SYNCH
    jz      rcvr_done         ; so drop this char
    cmp     ah,RS_STATE_1     ; we have been waiting for
40    jz      synch_1          ; this SYNCH to come along
                                ; to start off a block

    cmp     ah,RS_STATE_3     ; we have already received
    jz      synch_3           ; the ADDRESS, so this
                                ; SYNCH makes a short
45    ; packet. discard it

    cmp     ah,RS_STATE_4     ; we have already received
    jz      synch_4           ; the ADDRESS and the
                                ; CONTROL, but no protocol.
                                ; therefore, dump this
50    ; packet
    cmp     ah,RS_STATE_5     ; we have received both the
    jz      synch_5           ; ADDRESS and the CONTROL
                                ; field and only the first
                                ; byte of the protocol
                                ; field.
55    ; this SYNCH therefore makes
                                ; for a short packet, so

```

```

; we will discard it
; Okay, here's the tough one
; so listen up : In order
; to get to state 7, a
; valid packet came in and
; we sent the message off
; to pSCS, but when we
; tried to get a NEW packet
; buffer for the NEXT
; packet,
; there were none left. The
; fact that a packet was
; just sent out means that
; we DID RECEIVE the SYNCH
; char, and so if we receive
; ANOTHER one it means that
; we have stayed in state 7
; for the duration of an
; entire packet. the upper
; layer needs to be made
; aware of the fact that
; a bad packet was received
; so that VJ compression
; will work correctly.

    cmp     ah,RS_STATE_7
    jz      synch_7

5
    pushd   0
    bad_packet:
    pop     esi
    mov     ecx,PPP_ERROR
    mov     edx,ppp_com3_interface
    call    ppp_q_write
    mov     ah,RS_STATE_2
    jmp     rcvr_done
; reset everything

30
    synch_1:
    mov     ah,RS_STATE_2
    jmp     rcvr_done
; everything's great so
; go to the next state

    synch_3:
    synch_4:
    synch_5:
35
    mov     ah,RS_STATE_2
    jmp     rcvr_done
; inc short_packet_count
; Drop this packet, and
; then check for more chars

    synch_7:
40
    mov     ecx,PPP_ERROR
    mov     edx,ppp_com3_interface
    mov     esi,4
    call    ppp_q_write
    mov     ah,RS_STATE_7
    jmp     rx_state_7
; stay in state 7 and try
; to get another buffer

45
    synch_6:
    sub     edi,ppp_com3_rx_buffer_start
    cmp     edi,2
50
; get length of packet
; it needs to at least have
; 2 bytes if we are using
; CRC16 FCS

    pushd   1
    jb      bad_packet
    add     esp,4
; short packet error
; send an error and reset

55
    cmp     bx,PPP_FCS16_GOOD_VALUE
    pushd   2
; does the CRC match ?
; CRC error

```

```

        jnl     bad_packet          ; no, so send error and discard
        add     esp,4

5       sub     edi,2                ; yes, but don't send CRC as
                                   ; part of packet
        mov     ax,word ptr ppp_com3_rx_protocol ; encode the
                                   ; protocol as the
        shl     eax,16               ; upper 16 bits of the
10      or      edi,eax              ; length field
        mov     ecx,PPP_PACKET_RECEIVED ; message_type
        mov     edx,ppp_com3_interface ; interface
        mov     esi,ppp_com3_rx_buffer_start ; message_ptr

15      call    ppp_q_write          ; This is the line of code that
                                   ; notifies upper layers that the
                                   ; PPP packet was indeed valid.

        or      eax,eax              ; did it fail ?
20      jz      drop_packet          ; yes, throw away the packet

        call    ppp_buffer_get_rcv   ; no, so get a new packet buffer
        mov     edi,eax
        inc     eax
25      or      eax,eax
        mov     ah,RS_STATE_7
        jz      rcvr_done
        mov     ppp_com3_rx_buffer_start,edi
drop_packet:
        mov     ah,RS_STATE_2
30      jmp     rcvr_done

```

#### Claims

- 35 1. A computer, comprising:
- a host central processing unit (CPU);  
an expansion bus operably coupled between said host CPU and a video controller; and  
40 a server controller connected to said expansion bus, wherein said server controller comprises a controller memory for storing display data forwarded from said host CPU to said video controller during a first time period prior to failure of said computer and during a second time period after reset of said computer.
- 45 2. The computer as recited in claim 1, wherein said expansion bus comprises an EISA bus.
3. The computer as recited in claim 1 or claim 2, wherein said video controller is coupled to a video memory such that both said video controller and said server controller are responsive to address information within a video address range for simultaneously storing said display data in both said video memory and controller memory, respectively.
- 50 4. The computer as recited in any of claims 1 to 3, wherein said host CPU is connected to a local CPU bus.
5. The computer as recited in claim 4, further comprising a bus interface unit coupled between said local CPU bus and said expansion bus.
- 55 6. The computer as recited in any of claims 1 to 5, wherein said server controller comprises detection logic coupled between said expansion bus and said controller memory for monitoring said display data forwarded across said expansion bus to said video controller and for thereafter writing said display data to said controller memory.

7. The computer as recited in claim 6, wherein said display data is forwarded across said expansion bus during times when said host CPU writes said display data to said video controller.
8. The computer as recited in any of claims 1 to 7, wherein said controller memory comprises:
- a local frame buffer for storing a current display screen of said display data forwarded to said video controller;
  - and
  - a temporary snapshot buffer coupled to receive the current display screen of said display data forwarded from said local frame buffer;
  - a snapshot buffer coupled to receive a previous display screen of said display data forwarded from said local frame buffer; and
  - comparison logic coupled between said temporary snapshot buffer and said snapshot buffer for determining a change between said previous display screen and said current display screen of said display data.
9. The computer as recited in claim 8, wherein said controller memory further comprises a current reset sequence buffer for storing said change during said second time period.
10. The computer as recited in claim 8, wherein said controller memory further comprises a current sequence buffer operably coupled to a current start buffer, and wherein said change is forwarded to said current sequence buffer and any overflow therefrom is forwarded to said current start buffer during said first time period.
11. A server controller, comprising:
- detection logic adapted for coupling to a server to determine the presence of display data forwarded from the server;
  - a controller memory coupled to said detection logic, wherein said controller memory comprises a local frame buffer and a current reset sequence buffer; and
  - said local frame buffer is configured to store a current display screen of said display data, and said current reset sequence buffer is configured to store a change between the current display screen of said display data and a previous display screen of said display data during a time period after which said server receives a first reset.
12. The server controller as recited in claim 11, wherein said change between the current display screen of said display data and the previous display screen of said display data is represented as a packet of digital data.
13. The server controller as recited in claim 12, wherein said packet of digital data comprises fewer binary values than a set of binary values used in storing either the current display screen or the previous display screen.
14. The server controller as recited in claim 12, wherein said packet of digital data comprises a byte of data indicative of a character to be displayed and another byte of data indicative of attributes given to the character to be displayed, said attributes are selected from the group comprising colour and intensity.
15. The server controller as recited in claim 12, wherein said packets are arranged in series representing a sequence of video screen changes which occur after said server receives the first reset.
16. The server controller as recited in any of claims 11 to 15, wherein said current reset sequence buffer is adapted for access by a terminal located remote from said server to display upon said terminal a sequence of video screen changes which occur after said server receives the first reset.
17. The server controller as recited in any of claims 11 to 16, further comprising:
- a previous reset sequence buffer configured to store said change between the current display screen of said display data and the previous display screen of said display data upon the server receiving a second reset prior to the first reset; and
  - said previous reset sequence buffer is adapted for access by a terminal located remote from said server to display upon said terminal a sequence of video screen changes which occur after the server receives said second reset, wherein the current reset sequence buffer is adapted for access by said terminal to display upon said terminal a sequence of video screen changes which occur after said server receives the first reset.

18. A server controller, comprising:

detection logic adapted for coupling to a server to determine the presence of display data forwarded from the server;

a controller memory coupled to said detection logic, wherein said controller memory comprises a local frame buffer and a current sequence buffer; and

said local frame buffer is configured to store a current display screen of said display data, and said current sequence buffer is configured to store a change between the current display screen of said display data and a previous display screen of said display data during a time period before which said server fails.

19. The server controller as recited in claim 18, wherein said change between the current display screen of said digital data and the previous display screen of said digital data is represented as a packet of digital data.

20. The server controller as recited in claim 19, wherein said packet of digital data comprises fewer binary values than a set of binary values used in storing either the current display screen or the previous display screen.

21. The server controller as recited in claim 19, wherein said packet of digital data comprises a byte of data indicative of a character to be displayed and another byte of data indicative of attributes given to the character to be displayed, said attributes are selected from the group comprising colour and intensity.

22. The server controller as recited in claim 19, wherein said packets are arranged in series representing a sequence of video screen changes which occur before said server fails.

23. The server controller as recited in any of claims 19 to 22, wherein said current sequence buffer is adapted for access by a terminal located remote from said server to display upon said terminal a sequence of video screen changes which occur before said server fails.

24. A method for accessing a sequence of display screens forwarded from a server to a video controller, comprising:

loading into a local frame buffer a current display screen of display data forwarded from the server to the video controller;

comparing the current display screen of display data to a previous display screen of display data previously loaded into said local frame buffer to produce a screen change packet;

repeating the steps of loading and comparing for a sequence of display screens compiled as a plurality of screen change packets which occur during a first time period immediately prior to failure of said server and during a second time period immediately after reset of said server; and

assessing said plurality of screen change packets by a terminal located remote from said server to display upon the terminal said sequence of display screens.

25. The method as recited in claim 24, wherein said loading comprises snooping an expansion bus for addresses within a specific address range sent over the expansion bus from the server.

26. The method as recited in claim 24, wherein said accessing comprises remote dialling from the terminal into a modem coupled to a buffer configured to store said plurality of screen change packets.







FIG.2a

— REMOTE INSIGHT MAIN MENU —

1. ACCESS REMOTE CONSOLE
2. VIEW EVENT LOG
3. VIEW STATUS INFORMATION
4. VIEW SERVER RESET SEQUENCE
5. VIEW SERVER CRITICAL ERROR LOG
6. MANAGE LOGIN INFORMATION
7. MANAGE ALERT INFORMATION
8. RESET REMOTE SERVER
9. DISCONNECT

<TAB>=NEXT <ENTER>=SELECT <F1>=HELP

FIG.2b

— VIEW SERVER RESET SEQUENCE —

1. PREVIOUS RESET SEQUENCE REPLAY
2. CURRENT RESET SEQUENCE REPLAY
3. FAILURE SEQUENCE REPLAY

<TAB>=NEXT <ENTER>=SELECT <ESC>=BACK <F1>=HELP

FIG.2c

— SERVER RESET SEQUENCE PLAYBACK —

PRESS <SPACE> TO BEGIN THE SEQUENCE REPLAY.

DURING THE REPLAY:

PRESS <ESC> TO STOP THE PLAYBACK AND RETURN TO THE MENU.

PRESS <ENTER> TO PAUSE OR RESUME THE PLAYBACK.

PRESS <1-9> TO INCREASE OR DECREASE TO PLAYBACK SPEED.  
(1=SLOWEST, 9=FASTEST)

<ESC>=BACK





FIG. 6b




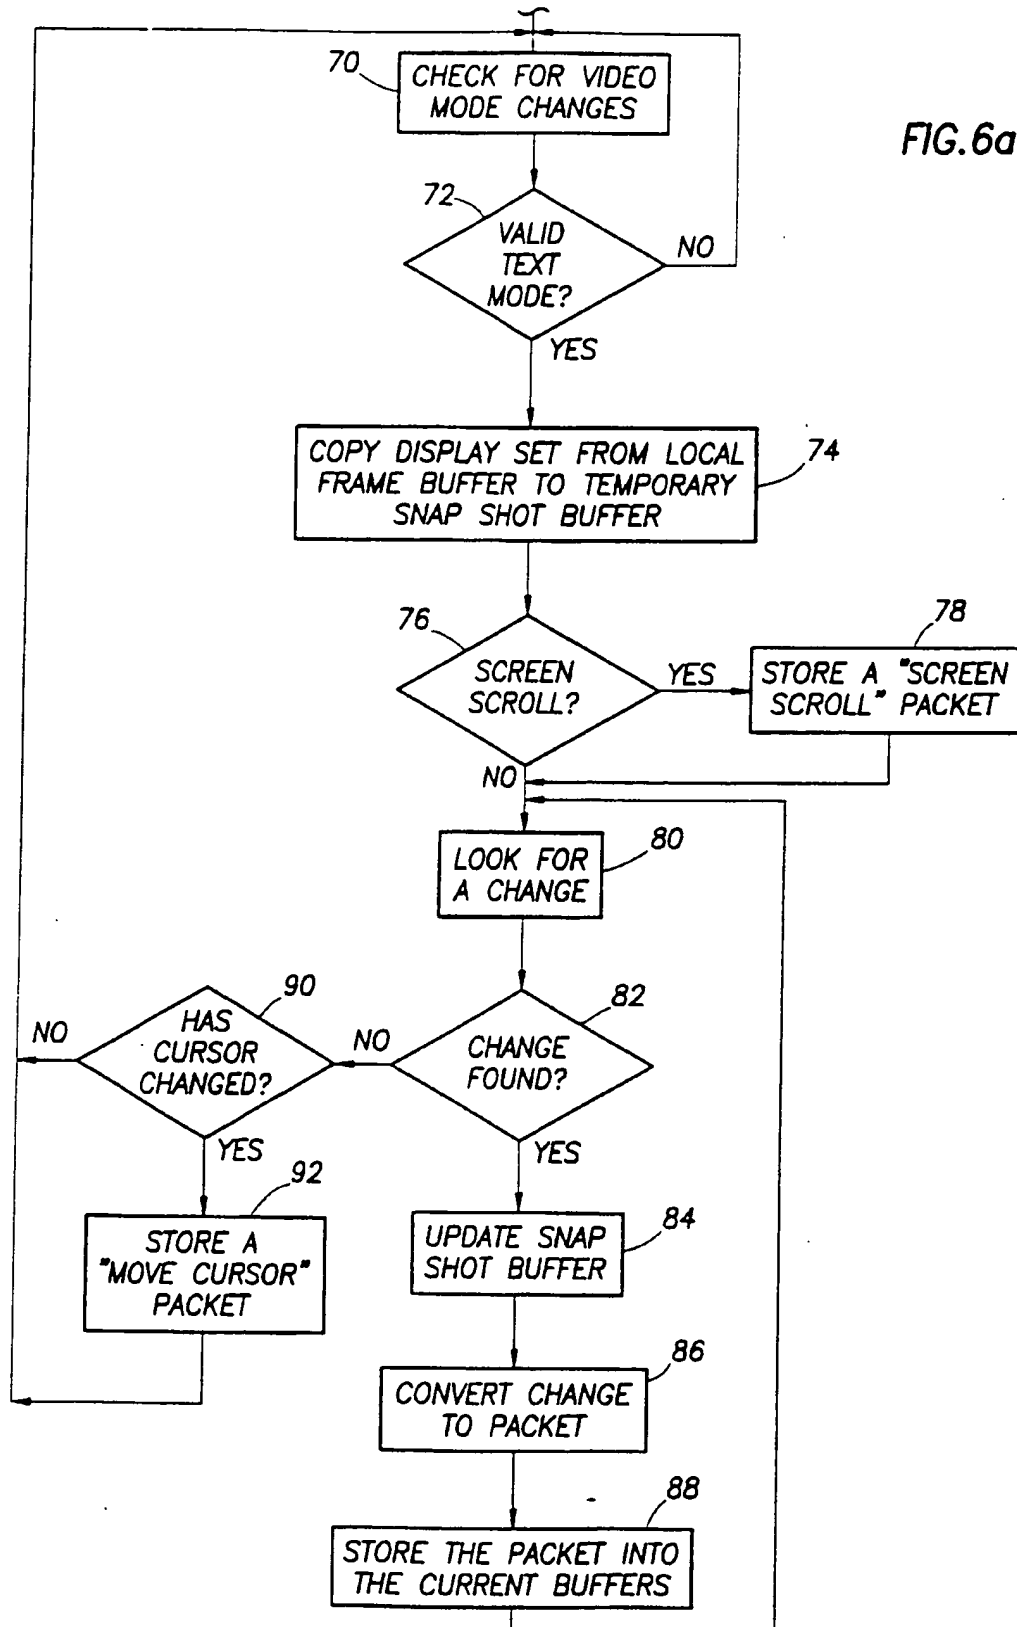
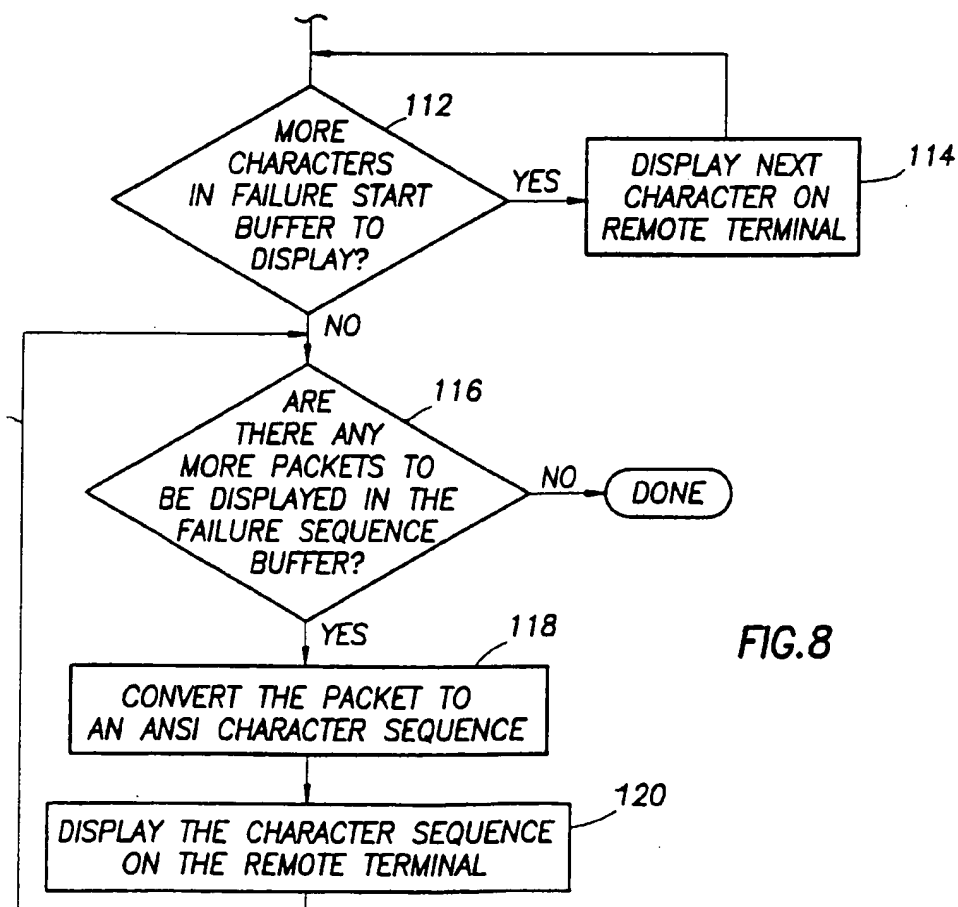



FIG. 6a





### FIG.9a

32768 KB DETECTED

COMPAQ SYSTEM BIOS-E15 (12/06/95)

COPYRIGHT 1982,1995 COMPAQ COMPUTER CORPORATION. ALL RIGHTS RESERVED.

PROCESSOR RUNNING AT 120 MHz

SCSI DEVICE DETECTED-SYSTEM BOARD, SCSI ID 0:COMPAQ M1606S

SCSI DEVICE DETECTED-SYSTEM BOARD, SCSI ID 5:COMPAQ CD-ROM CR-503BCQ

### FIG.9b

OS LOADER V3.51

PLEASE SELECT THE OPERATING SYSTEM TO START:

WINDOWS NT SERVER VERSION 3.51 FROM COMPAQ

WINDOWS NT SERVER VERSION 3.51 FROM COMPAQ (DEBUGGER ENABLED)

WINDOWS NT SERVER VERSION 3.51 FROM COMPAQ (VGA MODE)

USE ^ AND v TO MOVE THE HIGHLIGHT TO YOUR CHOICE.  
PRESS ENTER TO CHOOSE.

SECONDS UNTIL HIGHLIGHTED CHOICE WILL BE STARTED AUTOMATICALLY: 11

### FIG.9c

MICROSOFT (R) WINDOWS NT (TM) VERSION 3.51 (BUILD 1057)

1 SYSTEM PROCESSOR [32576 Kb MEMORY]

COMPAQ UNIPROCESSOR HAL-VERSION 301.01 FROM SSD 1.18

### FIG.9d

MONITORED SERVER CONSOLE IS IN GRAPHICS MODE.

### FIG.9e

STOP 0X0000001e (c000009a 80123f36 02000000 00000246)

UNHANDLED KERNEL EXCEPTION c000009a FROM 8123f26

ADDRESS 80123f36 HAS BASE AT 80100000-NTOSKRNL.EXE

FIG.10






FIG. 11

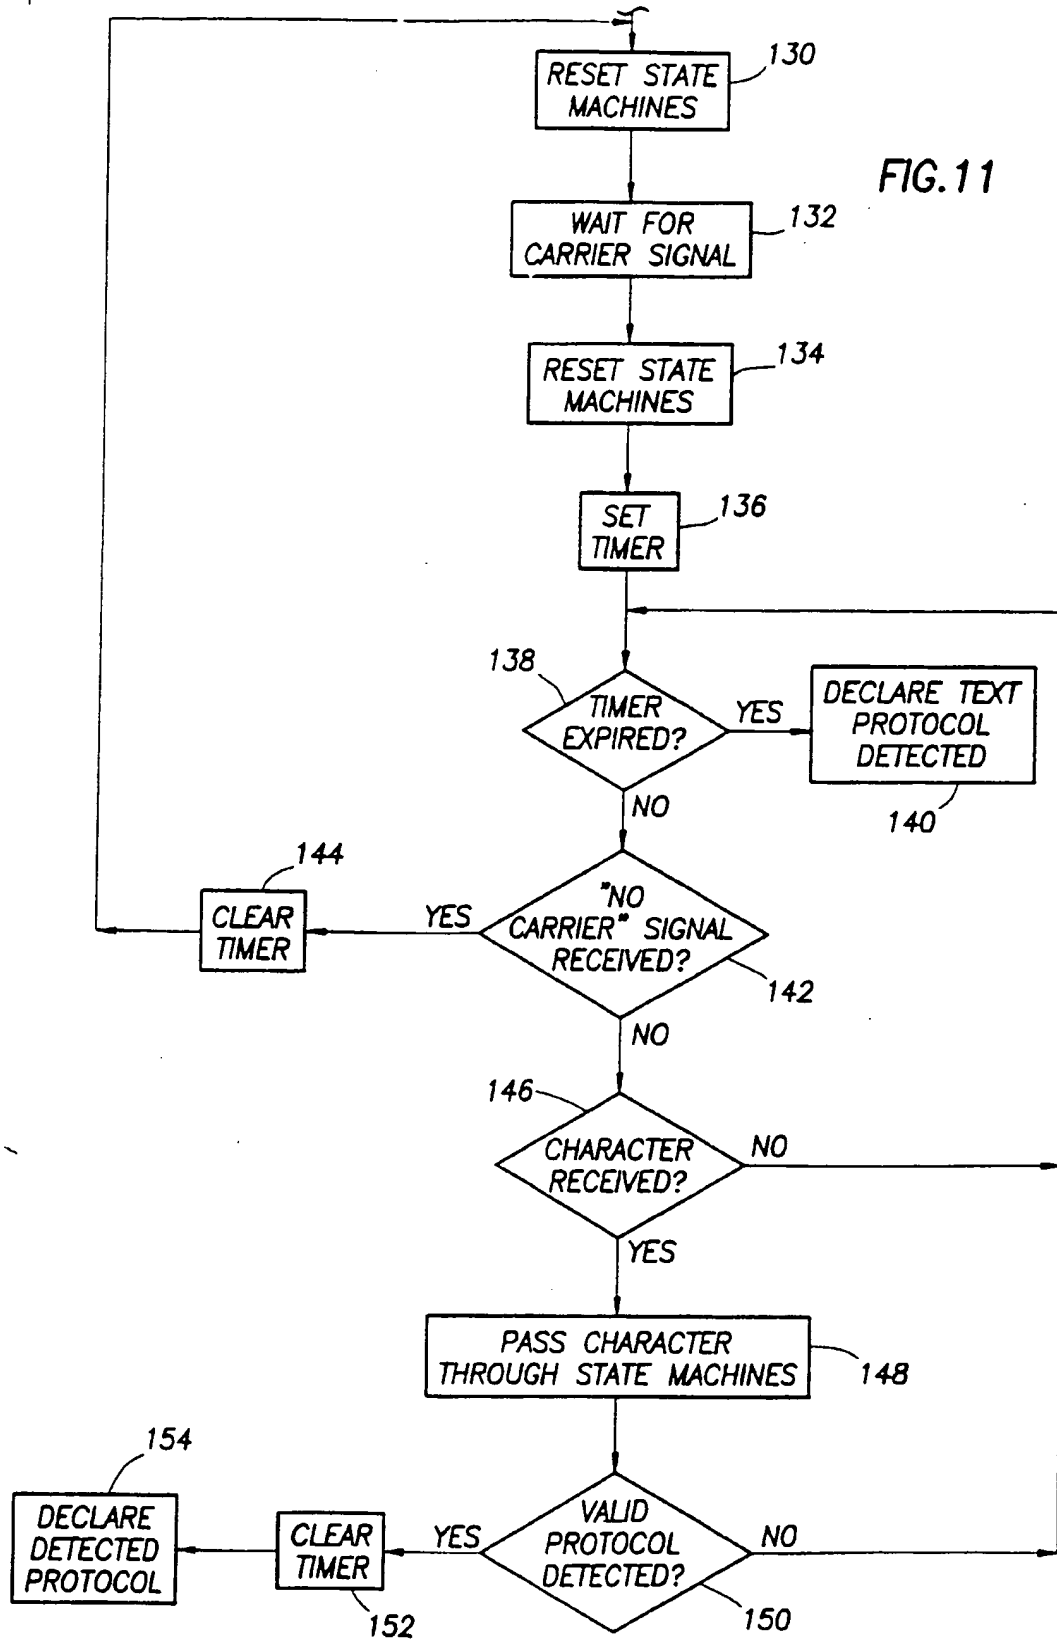


FIG. 12




FIG. 13

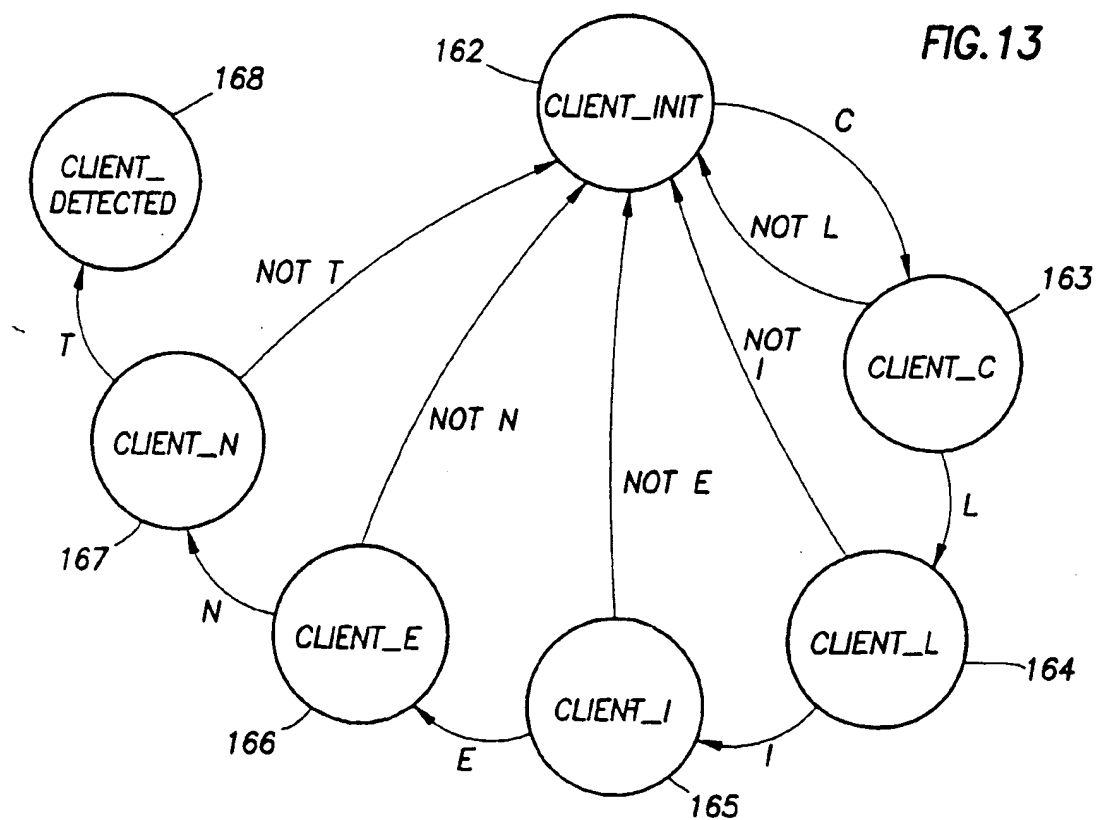


FIG. 14

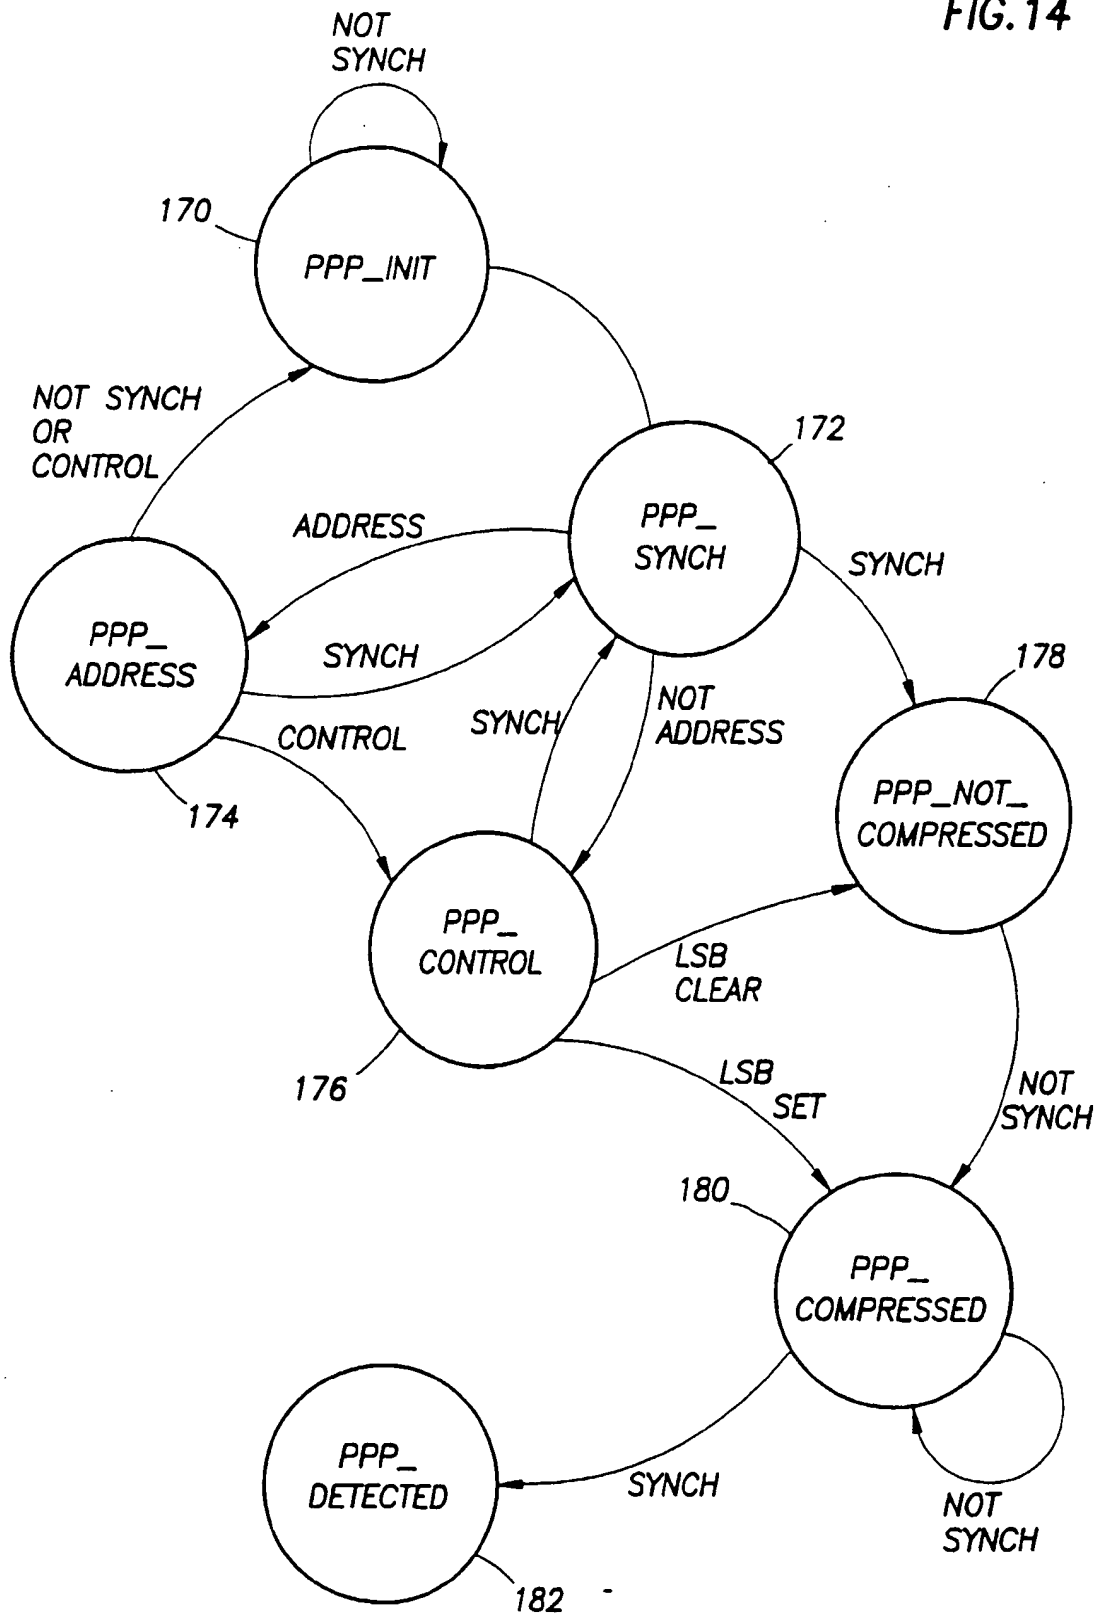


FIG. 15




FIG. 16






FIG. 17





(19)

Europäisches Patentamt

European Patent Office

Office européen des brevets



(11)

EP 0 825 535 A3

(12)

## EUROPEAN PATENT APPLICATION

(88) Date of publication A3:

10.03.1999 Bulletin 1999/10

(51) Int Cl.6: G06F 11/22

(43) Date of publication A2:

25.02.1998 Bulletin 1998/09

(21) Application number: 97306152.6

(22) Date of filing: 13.08.1997

(84) Designated Contracting States:

AT BE CH DE DK ES FI FR GB GR IE IT LI LU MC  
NL PT SE

Designated Extension States:

AL LT LV RO SI

(30) Priority: 16.08.1996 US 695828

(71) Applicant: Compaq Computer Corporation  
Houston Texas 77070 (US)

(72) Inventors:

• Gready, Scott R.  
Spring, Texas 77379 (US)

• Ellinger, Wesley M.


Houston, Texas 77064 (US)

• Clark, Gordon R.

The Woodlands, Texas 77380 (US)

(74) Representative: Brunner, Michael John et al  
GILL JENNINGS & EVERY  
Broadgate House  
7 Eldon Street  
London EC2M 7LH (GB)(54) **A server controller for storing and retrieving sequences of video screens forwarded by a server host computer**

(57) A communication system is presented whereby sequences of video screens sent from a host CPU to a video controller can be stored and subsequently retrieved by a terminal located remote from the host CPU. The host CPU and video controller form part of a server arranged within a distributed computing system. An administrator situated at the remote terminal can retrieve select video screens produced during server operations to determine information regarding the server configuration and possible causes of server failure or future failure. The sequence of video screens thereby represent video screen changes stored upon a server controller adapted for coupling to the server expansion bus. The video screen changes represent a sequence of video screen changes occurring prior to server failure or after server reset. Those changes provide beneficial information to an administrator located remote from the server, and allows the administrator to communicate with the server using several possible communication protocols. The server controller snoops display data written from the host CPU to the video controller and mirrors the display data upon buffers within the server controller. Information within the buffers can be called upon by a remotely situated administrator regardless of whether server power is lost in the interim.



EP 0 825 535 A3



European Patent  
Office

# EUROPEAN SEARCH REPORT

Application Number  
EP 97 30 6152

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.6)
P, X	"Remote Server Management With Integrated Remote Console" COMPAQ COMPUTER CORPORATION TECHNOLOGY BRIEF, 21 October 1996, pages 1-10, XP002089523 US Obtained from the Internet on January 4th 1999 at the URL: <ftp.compaq.com/pub/supportinformation/papers/582a1096_a4.pdf> * the whole document *	1-26	G06F11/22
A	EP 0 520 768 A (COMPAQ COMPUTER CORP) 30 December 1992 * the whole document *	1-26	
A	EP 0 520 766 A (COMPAQ COMPUTER CORP) 30 December 1992 * the whole document *	1-26	
			TECHNICAL FIELDS SEARCHED (Int.Cl.6)
			G06F H04L
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 8 January 1999	Examiner Fonderson, A
<p>CATEGORY OF CITED DOCUMENTS</p> <p>X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document</p> <p>T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons Δ : member of the same patent family, corresponding document</p>			

EPO FORM 1503 03/82 (P04C01)

**ANNEX TO THE EUROPEAN SEARCH REPORT  
ON EUROPEAN PATENT APPLICATION NO.**

EP 97 30 6152

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report.  
The members are as contained in the European Patent Office EDP file on  
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

08-01-1999

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP 0520768 A	30-12-1992	CA 2072198 A	25-12-1992
		US 5440699 A	08-08-1995
EP 0520766 A	30-12-1992	AT 160234 T	15-11-1997
		CA 2072178 A	25-12-1992
		DE 69223090 D	18-12-1997
		DE 69223090 T	07-05-1998
		US 5402431 A	28-03-1995

EPO FORM P0469

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82

**THIS PAGE BLANK (USPTO)**



**This Page is Inserted by IFW Indexing and Scanning  
Operations and is not part of the Official Record**

## **BEST AVAILABLE IMAGES**

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☐ **BLACK BORDERS**
- ☐ **IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- ☐ **FADED TEXT OR DRAWING**
- ☐ **BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- ☐ **SKEWED/SLANTED IMAGES**
- ☐ **COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- ☐ **GRAY SCALE DOCUMENTS**
- ☐ **LINES OR MARKS ON ORIGINAL DOCUMENT**
- ☐ **REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- ☐ **OTHER: \_\_\_\_\_**

**IMAGES ARE BEST AVAILABLE COPY.**

**As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.**

**THIS PAGE BLANK (USPTO)**